MIT Technology Review

Reviews p. 74

Should We Fear Artificial Intelligence?

Reviews p. 83

In Defense of Silicon Valley

Business Report p. 65

The Future of Money

MAR/APR 2015 $6.99

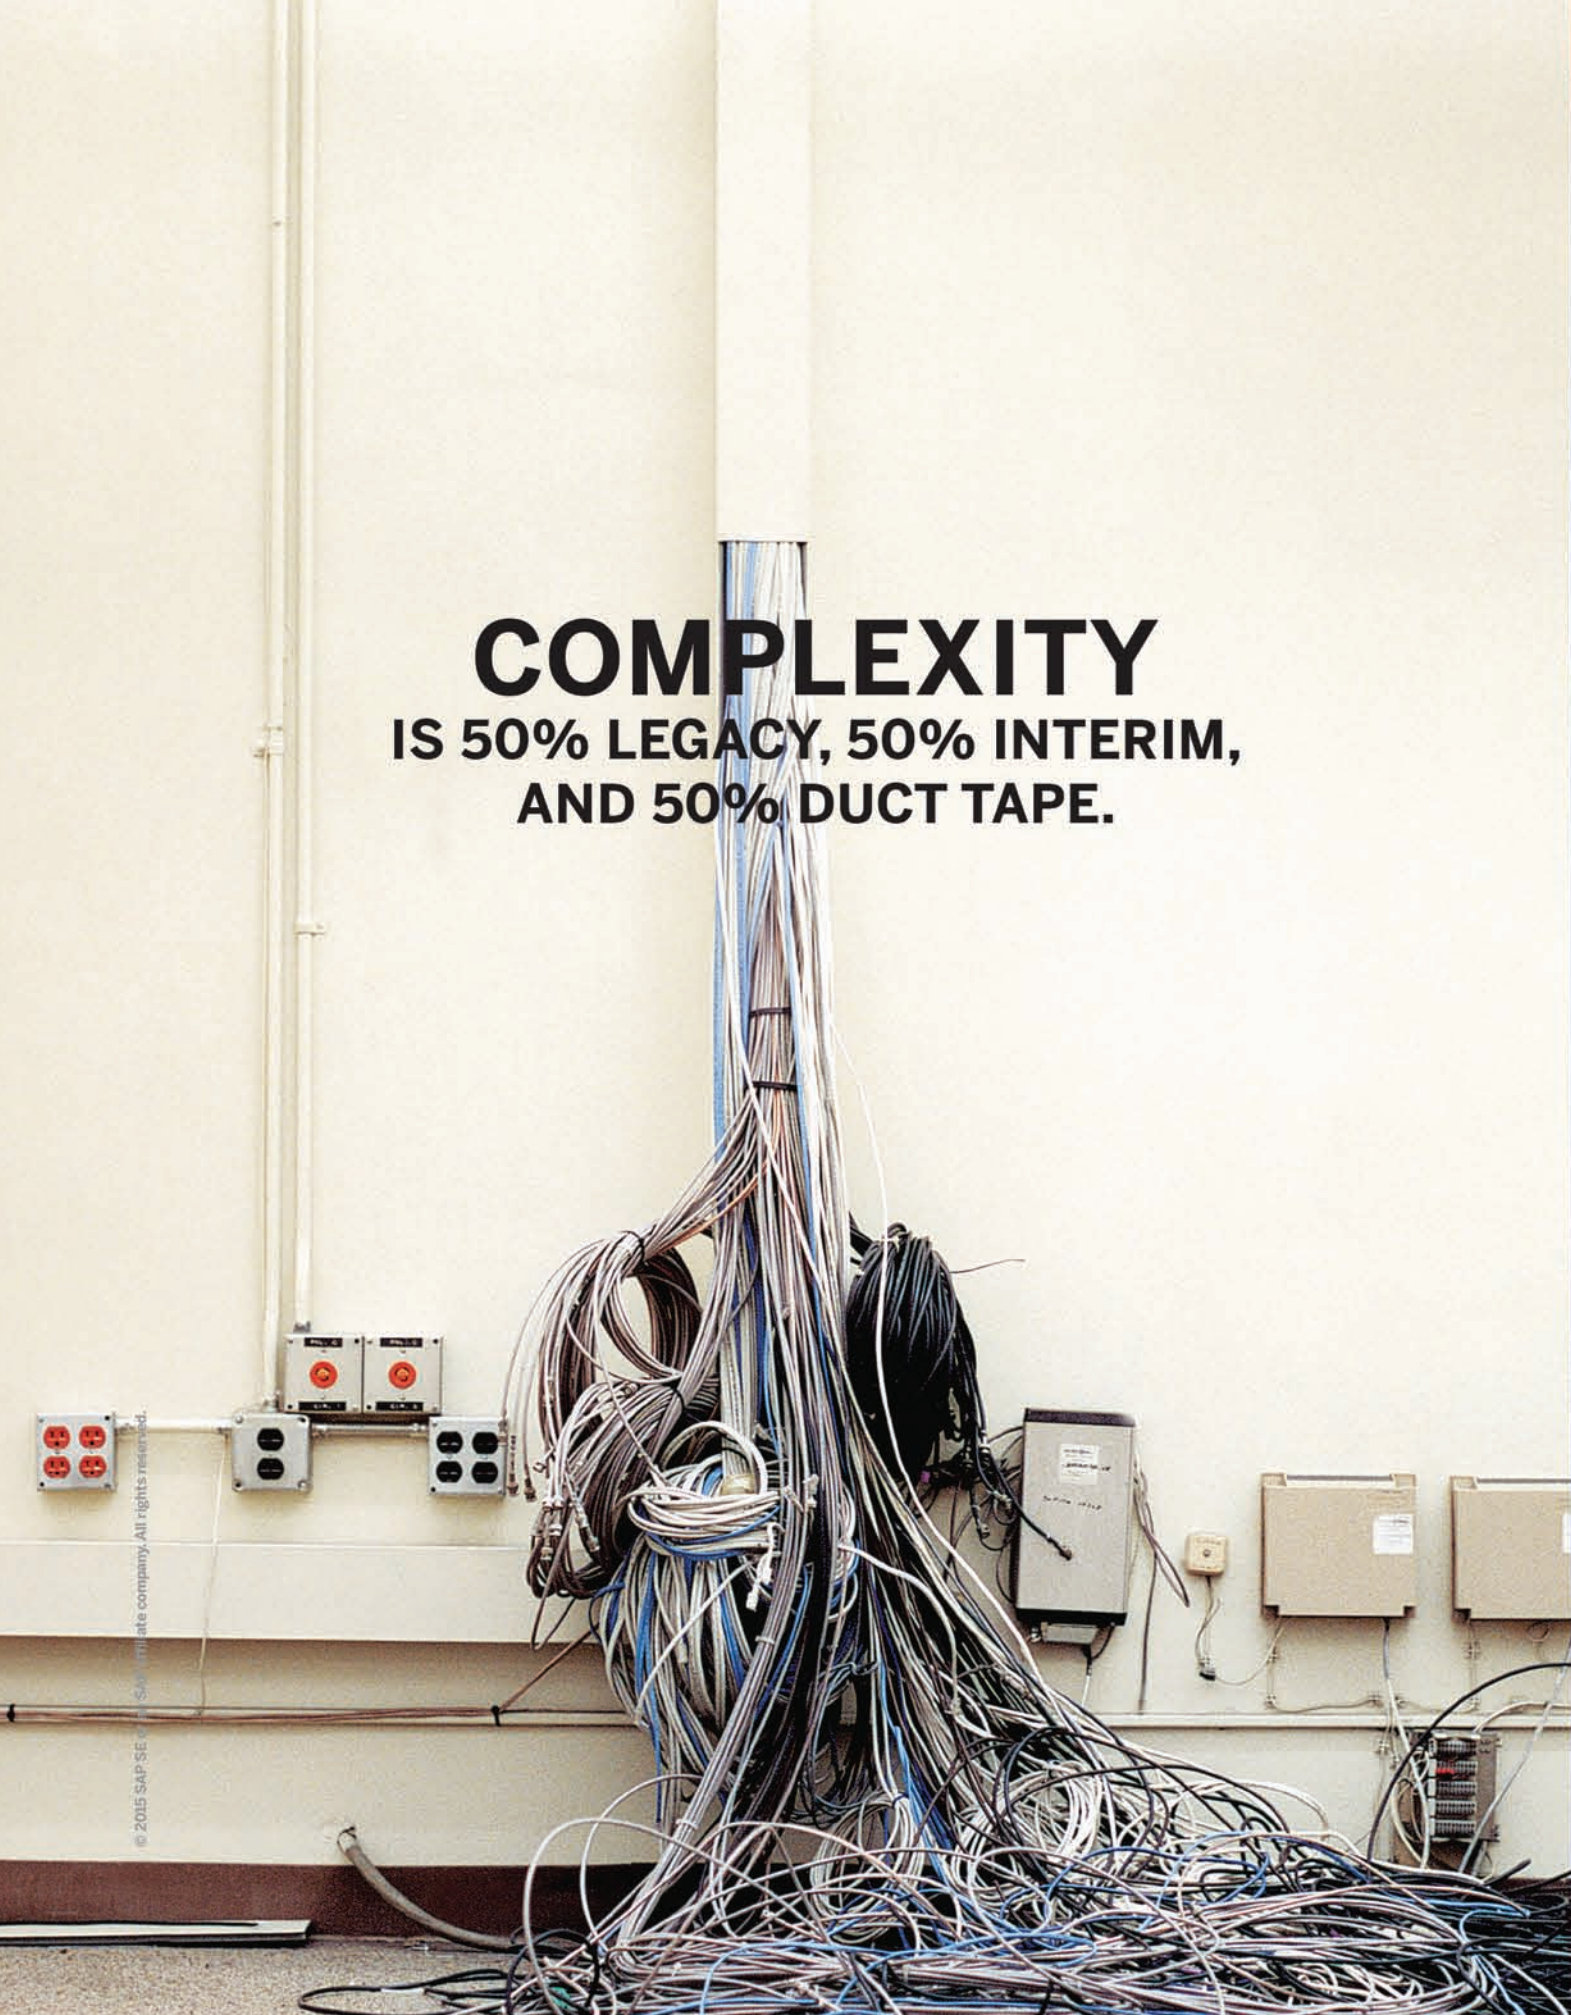


SIMPLE
JUST WORKS.
After decades of bolting one quick-fix technology on top of another, businesses now find themselves in a tangled IT environment that multiplies costs and stifles innovation. SAP puts all your applications on a single integrated innovation platform, agile enough to work with your current technology and innovate for the future. That's running simple. Find out more at sap.com/runsimple
SAP
Run Simple

From the Editor

GUIDO VITTI

EVERY YEAR, *MIT TECHNOLOGY REVIEW* selects the 10 technologies we believe are the greatest breakthroughs of previous months, those that in the future will have the broadest impact on commerce, medicine, and society.

The challenge and fascination of editing our publication (and therefore of creating this list) is that unlike many other technology magazines and websites, we are interested in all technologies, and most of all in how breakthroughs in one field may spur innovations in another. Those who have attended one of our seven EmTech events around the world may have watched as I struggled to explain how new developments in artificial intelligence (see "Deep Learning," one of our breakthrough technologies from 2013) may be connected to more efficient use of advanced renewable energy sources through predictive modeling (see "Smart Wind and Solar Power," a breakthrough technology from our 2014 list).

Our predictions are not always right, but even when we're wrong, we're interestingly mistaken. A few years ago, we correctly intuited that social media would be important to television (see "Social TV," a breakthrough technology in 2010). But we didn't understand that social and broadcast media wouldn't merge on TV screens; instead, people would watch television and update their impressions on Twitter, Facebook, or Instagram using their smartphones.

More commonly, we're not so much wrong as simply early: cancer genomics, where gene sequencing identifies the mutations behind an individual patient's specific cancer in order to more precisely identify the drugs most likely to work, was less practicable when it cost $30,000 to sequence someone's cancerous and healthy tissue (see "Cancer Genomics," from our 2011 list). Today, when the cost of sequencing a genome can be as low as $1,000, the president of the United States can talk at the State of the Union address about "precision medicine" as an imminent clinical reality. No matter that we jumped the gun a little: we prefer to be early than late.

This year, the 10 breakthrough technologies are similarly broad in scope. Senior editor Tom Simonite describes Google's Project Loon (page 40), an ambitious experiment by the company's Google X division to bring Internet access to the 60 percent of the world that doesn't have it by floating an armada of balloons with solar-powered electronics in the upper atmosphere.

Elsewhere, we report on cerebral organoids, clumps of tissue that possess certain features of the brain, which could "open a new window into how neurons grow and function, and … change our understanding of everything from basic brain activities to the causes of schizophrenia and autism" (see "Brain Organoids," by Russ Juskalian, on page 54).

Or consider the consumer technology Apple Pay (page 50): Robert Hof writes, "None of the individual technologies in it is novel, but the extent of Apple's control over both the software and the hardware in the iPhone—which exceeds what Google can do for Google Wallet even on Android phones—allowed it to combine those technologies into a service demonstrably easier to use than any other." Hof argues that Apple Pay will probably succeed in making mobile payments broadly used, where previous attempts have failed.

But read about all 10 technologies and write to tell me what you think at jason.pontin@technologyreview.com.

Contents

Front

MARCH/APRIL 2015

10 Breakthrough Technologies 2015

They'll change the way we drive, the way we eat, the way we cure disease. They'll change the way we access the Internet and provide drinkable water. Here are the advances most likely to change the way we see the world in the coming decades.

Back

ON THE COVER:

Illustration by Elliott Earls

Stop the Pirates Who Are Parroting Your New Ideas

Allen, Dyer, Doppelt, Milbrath & Gilchrist, P.A., an intellectual property law firm established in 1972, has protected the intrinsic rights of clients' original ideas and new technologies through the application of patents, trademarks,copyrights, licensing, infringement protection, trade secrets protection and unfair competition litigation for 42 years.

ADDM&G has garnered many awards and acknowledgements, including currently being listed as one of the Top 50 Law Firms in the nation, as recognized by *Trademark Insider,* as one of the Top 100 Patent Law Firms in the nation, as recognized by *Intellectual Property Today,* top listed in *Best Lawyers* magazine and in the Best Law Firms list in *U.S. News and World Report.* The firm has six attorneys with Florida Bar Board Certifications in IP Law and one attorney Board Certified in Business Litigation.

ADDM&G serves clients locally, statewide, nationally and internationally. Protect your new brainstorm with the brain trust of an Intellectual Property law firm. Contact us today for further information.

ADDM&G

ALLEN, DYER, DOPPELT, MILBRATH & GILCHRIST, P.A.

INTELLECTUAL PROPERTY ATTORNEYS

Orlando · 407·841·2330 | Jacksonville* · 904·398·7000 | Miami · 305·374·8303
Tampa* · 813·639·4222 | Winter Springs · 407·796·5064 | www.addng,com

GRAY MATTER MATTERS

*Satellite Office

CUSTOMER SERVICE AND SUBSCRIPTION INQUIRIES
National: 800-877-5230

International: 903-636-1115

E-mail: technologyreview@pubservice.com

Web: www.technologyreview.com/customerservice

MIT Records: 617-253-8270 (alums only)

Reprints: Donna Summers
dsummers@wrightsmedia.com
281-419-5725

Licensing and permissions:
licensing@technologyreview.com

Technology Review
One Main Street, 13th Floor
Cambridge, MA 02142
Tel: 617-475-8000

The mission of *MIT Technology Review* is to equip its audiences with the intelligence to understand a world shaped by technology.

Technology Review, Inc., is an independent nonprofit 501(c)(3) corporation wholly owned by MIT; the views expressed in our publications and at our events are not always shared by the Institute.

De Technologia non multum scimus. Scimus autem, quid nobis placeat.

MIT Technology Review

EmTech DIGITAL

MIT Technology Review Subscribers
SAVE UP TO $724
Use code ***TRSub*** at registration.

Innovations and Ideas Fueling Our Connected World

Speakers include:

Stewart Butterfield
CEO, Slack
Rethinking Digital Collaboration at Work

Mike Cassidy
Vice President, Google
Expanding the Reach of Wireless Internet Around the Globe

Adam Cheyer
Cofounder and VP of Engineering, Viv Labs
A New Era for Intelligent Digital Interfaces

John Maeda
Design Partner, Kleiner Perkins Caufield & Byers
Technology and Design in the Digital Era

Rosalind Picard
Chief Scientist, Empatica
Developing groundbreaking wearable devices with medical quality sensing

June 1–2, 2015 San Francisco, CA St. Regis Hotel

EmTech Digital explores tomorrow's digital technologies and explains their global impact on business and society. You'll gain access to the innovative people and companies at the heart of the next wave of the digital revolution.

Digital technologies are connecting everything around us, making everyday things smarter and augmenting our experiences. Information is the new utility. We'll look at emerging opportunities in the digital economy, and the implications for our personal and professional lives.

Featured topics include:

- Digital Life
- Digital Work
- Digital Play

Questions? Please e-mail:
eventsreg@technologyreview.com

REGISTER BY MARCH 15 AND USE CODE TRSub TO SAVE UP TO $724:
technologyreview.com/EmTechDigital

Feedback

E-mail letters@technologyreview.com

Write *MIT Technology Review*
One Main Street, 13th Floor
Cambridge, MA 02142

Please include your address, telephone number, and e-mail address.

Letters and comments may be edited for both clarity and length.

Five Most Popular Stories

MIT Technology Review
Volume 118, Number 1

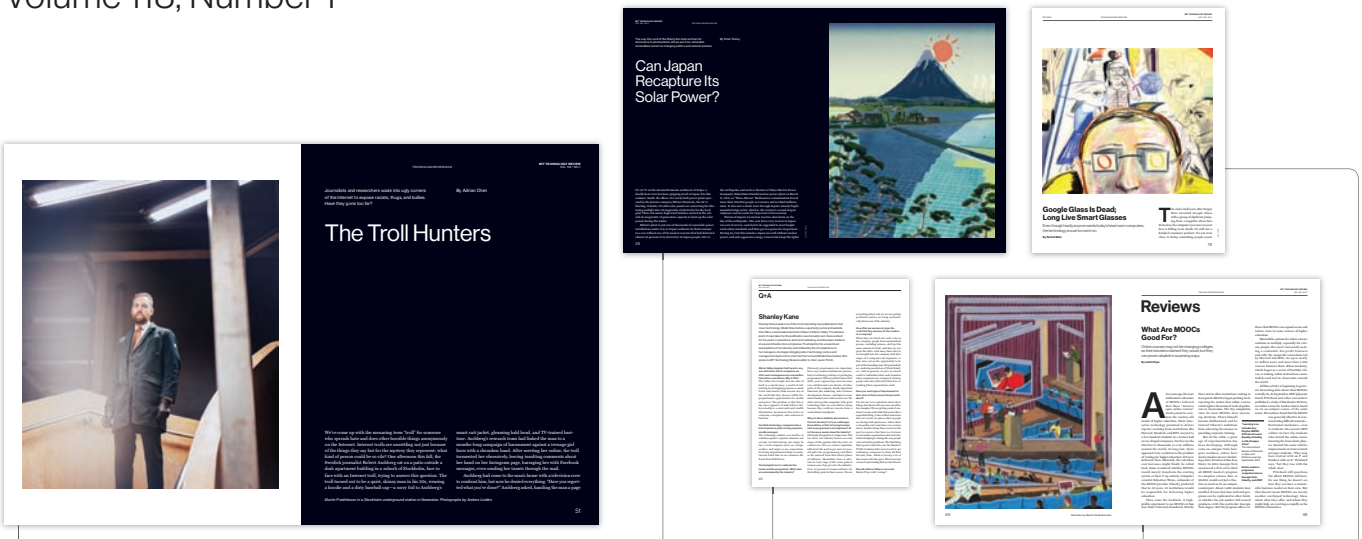

1 The Troll Hunters

A group using technology to expose hatred rather than reinforce it—I love it. If elected officials are spreading racist or xenophobic material online, the voting public has a right to know.

—thisoceanhasteeth

The entire idea of trolling and troll hunting is pure BS. It's just an excuse to silence opposing voices. We need to allow people to voice their opinion freely without being censored. We can think it's wrong for skinheads to spout their venom, but blocking them does not make their culture go away.

—William Ganness

2 Can Japan Recapture Its Solar Power?

Japan doesn't need solar power. It needs clean, abundant, reliable, and cheap energy. Nuclear power meets all these criteria. **—pobembe**

I'd suggest a follow-up called "Can Japan Rebuild Its Nuclear Power?" to analyze the lessons of Fukushima. What led to the faulty design? What safety reviews were held, and how were they approved? Where around the world have similar errors been made? Optimism can't substitute for science, engineering, and economics.

—David Korenstein

3 Q&A: Shanley Kane

Count me among the 50 percent of women who left the tech world. Despite 20 years of experience, at 40 I was presumed to be too old to understand the new ways. I was tired of the frat-boy mentality and outright sexism, and moved to academia. Much happier now. **—Bette Page**

The barrage of negativity about how "hostile" the industry is does nothing but discourage women from pursuing tech careers. I spent most of my 15 years in the industry under female managers. It's a truism—those that can, do; those that can't, whine. **—GlugGlugGlug**

4 What Are MOOCs Good For?

I took a music MOOC from Berklee and was touched by [something shared by] an Iraqi girl who was not allowed access to instruments. She was grateful for the MOOC but had to remain anonymous.

—Mike Reding

When I read that of 17,000 students in one MOOC, only 5 percent completed it, I view that with optimism. That's 850 students who completed the course. That's pretty impressive, given that the certificates are virtually worthless and most students were taking the course simply to learn.

—Richard Spears

5 Google Glass Is Dead

For those with peripheral vision loss, Glass literally gives sight to the blind. Thanks to Glass, my son sees and experiences life more fully. To dismiss a device before its use is fully explored is folly. That's exactly why it's an "explorer" program.

—embrownconroy

Google mismanaged the story of what Glass could be. It should have been positioned as a product for surgeons, mechanics, etc. Then the whole privacy debate would not even have arisen. But that's not Google's style. They like moon shots. **—Tim Meyer**

Humans = A Work in Progress

It's clear from Stephen S. Hall's article "Solving the Autism Puzzle" (January/February 2015) that our genome is not set up to fail. Rather, it is set up to change. We are a work in progress. Our genome has not yet become stable, if it ever will. Hall's article makes clear that we don't know where our genome is headed. Now that we are getting so smart, maybe we will attempt to guide the result. Or maybe the notion that we are smart enough to do that right is pure hubris. **—Allan B. Wheeler**

An Important Difference

Regarding John Elder Robison's article "Fixing Autism Research" (Views, January/February 2015): it bears repeating that not all of us [with autism] are affected in the same ways. For example, I was linguistically precocious and very bright but adversely affected in other ways, and it is taking a very long time to achieve near-normal function socially and in other ways—and I'm 70.

—ka5s

"Coping" Doesn't Get Grants

Robison is probably right to suggest that autism research needs to move from "cure" to "care" for well-being, but I think we are dreaming if we think that will actually happen. I'm a music therapist, and every time I've tried to obtain a grant for the application of music-based treatment approaches for autism spectrum disorder, I've been rejected because I wasn't discussing cells and molecules, only "coping" and "stress and anxiety reduction."

—dorita.berger

Robison is right to suggest that autism research needs to move from "cure" to "care" for well-being, but we're dreaming if we think that will actually happen.

Race for the Cure

I think John Elder Robison is a terrific person and I have enjoyed his books, but in this case I think it would be a mistake to follow his suggestion to stop looking for a cure.

That would simply lead to a longer time before any meaningful help can be offered to those with autism. Yes, autism is a number of diseases, but there is a dominant phenotype that is biochemical in nature. What we need is researchers who will delve into the biochemistry of this condition. Even if you think the search for "cures" is misguided, based on health reasons alone it makes sense to try to address the imbalances of biochemistry that characterize most cases of autism. Many people with it suffer from conditions including seizures, asthma, eczema, autoimmune disease, and Parkinson's as they get older. Why is this, exactly? It is because of dysfunctional biochemistry, and we need to do the research so we can understand it better. **—Seth Bittker**

Put a Price on Water

David Talbot writes that desalination is "one of the most expensive sources of fresh water," costing between $1,000 and $2,500 to supply two five-person households ("Desalination Out of Desperation," January/February 2015). Calling this is "expensive" is nonsense. I pay around $800 per year, and I live in the Great Lakes region. The biggest impediment to widespread deployment of desalination is the unwillingness to put a price on water. Do so—with a subsidy if you like for low-income families—and water shortages will cease. **—pasward**

Views

Seval Oz

Kosta Grammatis

Sharon Clay

TRANSPORTATION

Clever Vehicles

The way we go online in our cars makes no sense.

MOST OF US ARE NOW USED TO HAVING the Internet and the cloud with us as we drive. The problem is that we're doing it the wrong way. We "import" the Internet into the car via our smartphones, which is not just inefficient but dangerous: it takes our focus off the road and overloads us with distracting information.

Instead, we should be building sensors and networking technology right into the car—essentially making the vehicle itself our way of interacting with the online world (see "Car-to-Car Communication," page 38). This kind of in-car system is becoming more and more common. It gathers data about the environment around the car, allowing vehicles to sense, know, and predict.

The safety implications of such systems are massive. Our cell phones are a distraction when we should be concentrating on the road. Yes, we all want to be connected all the time, and the time we spend in the car shouldn't be any exception—we just have to be smart about it. You don't want to be taking a sales call when you're making a left-hand turn at an intersection. An intelligent system would enable the car to prioritize data about the car itself and about the environment around the car, delivering only the information the driver needs at the exact moment it's needed.

A car equipped with such technology could initiate the appropriate communication when, say, a driver is returning home and wishes to turn up the thermostat there. The human driver would not have to think about it. The car would be able to fuse real-time data from a controller area network with mapping and timing data to "predict" that the driver was heading home.

A smart system wouldn't just remove distractions—it could help us avoid dangerous conditions that we can't see. This will be increasingly true as more vehicles are equipped with longer-range sensors and the ability to communicate with each other. Cars on the street will form networks—thanks to other cars on the road, your car can tell you about the road conditions you're about to encounter, accidents that you can avoid or just can't see, or congestion that's just in the process of developing.

We already live in an online world. It's only natural that our cars will become part of that world. A vehicle should be no different from an office, a home, or an airport. The technology to establish this consistent data flow in, around, and beyond the car is ready, and it's about time we started taking advantage of it.

Seval Oz is CEO of Continental Intelligent Transportation Systems and the former head of business development for Google's self-driving-car program.

INTERNET

The Web as a Right

Using the Internet isn't just a luxury.

WHAT WE'VE BEEN CALLING THE WORLD Wide Web for the last few decades isn't actually worldwide. Four billion people, more than half the world's population, remain disconnected. No e-mail, no Wikipedia. Six years ago I started a nonprofit called A Human Right to help connect the disconnected. At the time, the topic wasn't a mainstream issue, but now Mark Zuckerberg, Elon Musk, Google, and Richard Branson are all working to ensure that the Web indeed becomes worldwide (see "Project Loon," page 40).

ANDY FRIEDMAN

Why is access so important? The Internet is bigger than Wikipedia and e-mail. I believe Internet access should be considered a basic human right.

In 1948, the United Nations adopted a Declaration of Human Rights that covers basic liberties like the right not to be enslaved, the right to access education, and the right to a home and a job. Obviously, those framers couldn't have predicted the existence of the Internet or what it would mean for humanity. If the declaration were written today, it's likely that Internet access would be included.

This doesn't mean that governments or taxpayers should be expected to provide Internet access to everybody around the world. It does mean that we should protect access for those who have it, and encourage expansion to the people who don't. It means that a government that attempts to restrict, impede, or remove Internet access is violating its citizens' basic human rights.

The online world enables many things the U.N. considers human rights, such as education and freedom of expression. In Kabul, Afghanistan, hundreds of young Afghan women have been poisoned for attending school and brutalized for expressing themselves. But a small women-only cybercafé, funded partially by online donors, has opened. It has, in effect, restored their rights to education and expression.

But Internet access isn't important just because it supports the exercise of other human rights—framing Web access that way undervalues the power of being connected.

The Internet, built on top of the brick-and-mortar society we call civilization, is its own unique society. It grants a global perspective to our lives. Internet access, in fact, makes people a part of a global citizenship—it gives them the ability to collaborate, learn, interact, and empathize with the rest of humanity.

It is the only place where all people can meet as equals to engage in the global digital society that shapes our physical world. Without access, the disconnected don't have a voice in the process, and the world moves on without them. With it, they can have an impact that's worldwide.

Kosta Grammatis is a former engineer at SpaceX and is now a correspondent for Al Jazeera covering technology.

BIOTECHNOLOGY

Pliable Plants

We need crops that can adapt to a less predictable environment.

PRODUCING HIGHER-YIELDING CROPS HAS been a goal for farmers ever since there have been farms. Plant breeders since Gregor Mendel have recognized the heritability of crop traits and how they relate to yield. Heritability is crucial, but it's only one factor—there's also the question of how stable an inherited trait will be under different environmental stresses.

We now have tools at our disposal that Mendel could never have dreamed of. The addition of genes for C4 photosynthesis into non-C4 plants, for example, gives us crops that are much more efficient at using water and nitrogen (see "Supercharged Photosynthesis," page 58).

But such advances are only worthwhile if they lead us to crops that yield well in the variable and unpredictable environments that we now face. In the northern Great Plains, as an example, farmers have seen temperatures increase and precipitation patterns change during critical pollination times. We need to develop plant varieties that can withstand these changes.

According to the U.S. Department of Agriculture, 73 million acres of farmland in the United States were taken out of production between 1990 and 2012. Similar trends can be seen worldwide. The reasons are many, but three broad categories are desertification, rising oceans, and urbanization.

This loss of land means that we need to use what we have left more efficiently. Take the average U.S. corn acre, which produced 171 bushels of corn in 2014, according to USDA estimates. A bushel of corn weighs about 56 pounds. If, say, 80 acres are urbanized, other areas need to increase production by 760,000 pounds to replace the lost food.

There are a variety of agronomic ways to do that, one being genetically modified crops. Many people object to such plants, but they've resulted in considerable improvements to human and environmental health. One major benefit is the increase in "no-tillage acres," which greatly reduces soil erosion. This in turn prevents silt from entering streams and reservoirs, prevents the runoff of agrochemicals, and increases the water-holding capacity of the land.

Consider what happened in South Dakota during droughts in 1972 and in 2012. In 2012 the corn yields were about 70 bushels per acre greater than they'd been in 1972, thanks in part to improved crop genetics and the soil's increased water capacity.

Agronomists continue to prepare for an uncertain future. Food insecurity is a real and tangible threat. Through genetic engineering, conventional breeding, and research into improved management practices, agronomists are developing crops with smaller carbon footprints while also increasing yield and nutritional values, and they're doing all this on a shrinking number of arable acres. No small feat.

Sharon Clay is a professor of plant science at South Dakota State University and a past president of the American Society of Agronomy.

PROFESSIONAL EDUCATION

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

YOUR DOOR TO MIT EXPERTISE & KNOWLEDGE

TRAINING & EDUCATION FOR PROFESSIONALS

Come to MIT for a Week

SHORT PROGRAMS

Register for a 2–5 day intensive course and gain critical knowledge to help advance your career and impact your company's success. Earn CEUs and a Certificate of Completion.

Each year approximately 40 courses are offered in a variety of subject areas (partial listing below):

- Biotechnology/ Pharmaceutical
- Crisis Management
- Data Modeling & Analysis
- Design, Analysis, & Manufacturing
- Energy/ Transportation
- Imaging
- Information & Communication Technologies
- Innovation
- Leadership
- Radar
- Robotics
- Sustainability
- Systems Engineering
- Tribology

Come to MIT for a Semester

ADVANCED STUDY PROGRAM

Enroll in MIT courses through this non-matriculating, non-degree program. Participate on a full or part-time basis for one or more semesters.

As an Advanced Study Program participant you can:

- Advance your knowledge in the latest technologies and cutting edge research
- Choose from over 2,000 MIT undergraduate and graduate courses
- Earn grades, MIT credit, and a Certificate of Completion

Bring MIT to You

ONLINE X PROGRAMS

Interact with professionals from around the world without disrupting your work-life balance. Online courses designed to advance your career at your own pace with a flexible schedule that meets the needs of a global workforce.

Tackling the Challenges of Big Data
6 week online course announced for Fall 2014 & Spring 2015. **HTTP://MITPROFESSIONALX.EDX.ORG**

INTERNATIONAL PROGRAMS

Schedule regionally relevant MIT short courses in a global location near you today. MIT Professional Education will work with you to match your regional needs with relevant subject areas in the Short Programs portfolio.

CUSTOM PROGRAMS

Enhance your organization's capabilities and expertise through customized programs tailored to meet your specific corporate education needs and strategic goals.

To learn more about what MIT Professional Education can offer you and your company, visit us today at **HTTP://PROFESSIONALEDUCATION.MIT.EDU/TECHREVIEW** or email us at **PROFESSIONALEDUCATION@MIT.EDU.**

Upfront

KEITH NEGLEY

Who Owns the Biggest Biotech Discovery of the Century?

There's a bitter fight over the patents for CRISPR, a breakthrough new form of DNA editing.

By Antonio Regalado

Upfront

Last November in Silicon Valley, biologists Jennifer Doudna and Emmanuelle Charpentier showed up in black gowns to receive the $3 million Breakthrough Prize, a glitzy award offered by Internet billionaires including Mark Zuckerberg. They'd won for developing CRISPR-Cas9, a "powerful and general technology" for editing genomes that's been hailed as a biotechnology breakthrough.

Not dressing up that night was Feng Zhang, a researcher in Cambridge at the MIT-Harvard Broad Institute. But earlier last year Zhang claimed his own reward. In April, he won a U.S. patent that could give him and his research center control over just about every important commercial use of the technology.

How did the high-profile prize for CRISPR and the patent on it end up in different hands? That's a question now at the center of a debate over who invented what, and when, that involves three heavily financed startup companies, a half-dozen universities, and thousands of pages of legal documents.

"The intellectual property in this space is pretty complex, to put it nicely," says Rodger Novak, a former pharmaceutical industry executive who is now CEO of CRISPR Therapeutics, a startup in Basel, Switzerland, that was cofounded by Charpentier. "Everyone knows there are conflicting claims."

At stake are the rights to an invention that may turn out to be the most important new genetic-engineering technique since the beginning of the biotechnology age in the 1970s.

The CRISPR system, dubbed a "search and replace function" for DNA, lets scientists easily disable genes or change their function by replacing DNA letters. During the last few months, scientists have shown that it's possible to use CRISPR to rid mice of muscular dystrophy, cure them of a rare liver disease, make human cells immune to HIV, and genetically modify monkeys. No CRISPR drug yet exists, but if CRISPR turns out to be as important as scientists hope, commercial control over the underlying technology could be worth billions.

The control of the patents is crucial to several startups that together quickly raised more than $80 million to turn CRISPR into cures for devastating diseases. This include Editas Medicine and Intellia Therapeutics, both of Cambridge, Massachusetts. Clinical trials could begin in as little as three years.

Zhang cofounded Editas Medicine, and in December the startup announced that it had licensed his patent from the Broad Institute. But Editas doesn't have CRISPR sewn up. That's because Doudna, a structural biologist at the University of California, Berkeley, was a cofounder of Editas too. And since Zhang's patent came out, she's broken off with the company, and her intellectual property—in the form of her own pending patent—has been licensed to Intellia, a competing startup. Making matters more complicated, Charpentier sold her own rights in the same patent application to CRISPR Therapeutics. In an e-mail, Doudna said she no longer has any involvement with Editas.

Few researchers are now willing to discuss the patent fight. But academic labs aren't waiting for the patent claims to get sorted out. Instead, they are racing to assemble very large engineering teams to perfect and improve the genome-editing technique.

The powerful new gene-editing system was unearthed in bacteria, which use it as a way to identify, and then carve up, the DNA of invading viruses. That work stretched across a decade. Then, in June 2012, a small team led by Doudna and Charpentier published a key paper showing how to turn that natural machinery into a "programmable" editing tool that let researchers cut any DNA strand, at least in a test tube.

The next step was clear—scientists needed to see if the editing magic could

TO MARKET

Smarter Cars

Drive PX and Drive CX

COMPANY:
Nvidia

PRICE:
Not disclosed

AVAILABILITY:
Now

Two new automotive computers will make it easier to build smarter vehicles. The Drive PX is capable of processing information from up to 12 cameras simultaneously, and it comes with software designed to assist with safety or autonomous driving systems. It also includes software trained to recognize different objects using a powerful technique known as deep learning. The Drive CX can power several high-resolution 3-D displays and perform natural-language processing to allow hands-free control. "It's pretty cool to bring this level of powerful computation into cars," says John Leonard, a professor of mechanical engineering at MIT, who has developed automated vehicles.

COURTESY NVIDIA

work on the genomes of human cells. In January 2013, the laboratories of Harvard's George Church and Broad's Zhang were first to publish papers showing that the answer was yes. Doudna published her own results a few weeks later.

Everyone by then realized that CRISPR might become an immensely flexible way to rewrite DNA, and possibly to treat rare genetic diseases. Venture capital groups quickly began trying to recruit the key scientists behind CRISPR, tie up the patents, and form startups.

No CRISPR drug yet exists. But if CRISPR turns out to be as important as scientists hope, commercial control over the underlying technology could be worth billions.

Charpentier threw in with CRISPR Therapeutics in Europe. Doudna had already started a small company, Caribou Biosciences, but in 2013 she joined Zhang and Church as a cofounder of Editas. With $43 million from leading venture funds Third Rock Ventures, Polaris Partners, and Flagship Ventures, Editas looked like the dream team of gene-editing startups.

In April of this year, Zhang and the Broad won the first of several sweeping patents that cover using CRISPR in eukaryotes, or any species whose cells contain a nucleus. That meant that they'd won the rights to use CRISPR in mice, pigs, cattle, humans—in essence, in every creature other than bacteria.

The patent came as a shock to some. Broad had paid extra to get it reviewed very quickly, in less than six months, and few knew it was coming. Along with the patent came more than 1,000 pages of documents. According to Zhang, Doudna's predictions in her own earlier patent application that her discovery would work in humans was "mere conjecture," while he was the first to show it, in a separate and "surprising" act of invention.

The patent documents have caused consternation. The scientific literature shows that several scientists managed to get CRISPR to work in human cells. In fact, its easy reproducibility in different organisms is the technology's most exciting hallmark. In patent terms, this could mean it was "obvious" that CRISPR would work in human cells, and that Zhang's invention might not be worthy of its own patent.

What's more, there's scientific credit at stake. In order to show he was "first to invent" the use of CRISPR in human cells, Zhang supplied snapshots of lab notebooks that he says show he had the system up and running in early 2012, even before Doudna and Charpentier published their results or filed their own patent application. That timeline would mean he hit on the CRISPR-Cas editing system independently. In an interview, Zhang affirmed that he'd made the discoveries on his own.

Not everyone is convinced. "All I can say is that we did it in my lab with Jennifer Doudna," says Charpentier, now a professor at the Helmholtz Centre for Infection Research and Hannover Medical School in Germany.

Lawyers for Doudna and Charpentier are expected to mount an interference proceeding in the U.S.—a winner-take-all legal process in which one inventor can take over another's patent. Who wins will depend on which scientist can produce lab notebooks, e-mails, or documents with the earliest dates.

"I am very confident that the future will clarify the situation," says Charpentier. "And I would like to believe the story is going to end up well."

PING ZHU

3 QUESTIONS

Megan Smith

You left Google to be the chief technology officer for the United States. Why does a country need a CTO?

The official role of the U.S. CTO is to advise the president and his team on how to unleash the power of data, innovation, and technology. You're not the VP of engineering. You are really an innovation instigator, very much like a commercial CTO often is.

What has surprised you most since joining the government?

What I love, I guess, and I sort of suspected from talking to people from the interview process, was that government is incredibly entrepreneurial ... it's actually full of entrepreneurs. And the other thing is, government has a lot of resources and if we get it aligned in the right way, we can really bring extraordinary service. We can make amazing things happen for the people of the nation if we come and work in these worlds.

How can you drive policy when Congress isn't technologically literate?

I think we are making great progress. The American people have a broad sense of opinions and perspectives, and Congress represents the American people, so we have to work with that just like we are working with the American people. I haven't been stopped from doing the important work that we are doing—getting technical folks in the government, getting open data. I have yet to meet anybody who has gone negative.

—Jason Pontin

Upfront

913

The number of "neurons" in a brainlike chip made by IBM using a powerful new memory material.

Google's Intelligence Designer

The man behind a startup Google acquired for $650 million plans to build a revolutionary new artificial intelligence.

By Tom Simonite

Demis Hassabis started playing chess at age four and soon blossomed into a child prodigy. At age eight, success on the chessboard led him to ponder two questions that have obsessed him ever since: first, how does the brain learn to master complex tasks; and second, could computers ever do the same?

Now 38, Hassabis puzzles over those questions for Google, having sold his little-known London-based startup, DeepMind Technologies, to the search company last year for a reported 400 million pounds ($650 million at the time). Google snapped up DeepMind shortly after it demonstrated software capable of teaching itself to play classic video games at a superhuman level. At the TED conference in Vancouver this year, Google CEO Larry Page gushed about Hassabis and called his company's technology "one of the most exciting things I've seen in a long time."

DeepMind's technology could improve some of Google's existing products. But if the technology progresses, it could change the role computers play in many fields. Building software that can learn when faced with almost any problem could help address some of the world's most intractable problems, says Hassabis. "AI has huge potential to be amazing for humanity," he says. "It will really accelerate progress in solving disease and all these things we're making relatively slow progress on at the moment."

Hassabis's quest to understand and create intelligence has led him through three careers: game developer, neuroscientist, and now artificial-intelligence entrepreneur. At 17, he led development of the classic simulation game Theme Park, released in 1994. He went on to complete a degree in computer science at the University of Cambridge and founded his own successful games company in 1998. But the demands of building successful computer games limited the time Hassabis could devote to his true calling. "I thought it was time to do something that focused on intelligence," he says.

So in 2005, Hassabis began a PhD in neuroscience at University College London, with the idea that studying real brains might turn up clues that could help with artificial intelligence. He chose to study the hippocampus, a part of the brain that underpins memory and spatial navigation, which is still relatively poorly understood.

He soon made a mark. In a 2007 study recognized by the journal *Science* as a "Breakthrough of the Year," he showed that five patients suffering amnesia due to damage to the hippocampus struggled to imagine future events. It suggested that a part of the brain thought to be concerned only with the past is also crucial to planning for the future.

That memory and forward planning are intertwined was one idea Hassabis took with him into his next venture. In 2011, he quit life as a postdoctoral researcher to found DeepMind, whose stated goal was to "solve intelligence." The company hired leading researchers in machine learning and attracted noteworthy investors, including Peter Thiel's firm Founders Fund and Elon Musk, founder of Tesla and SpaceX. But DeepMind kept a low profile until December

MARC ASPINALL

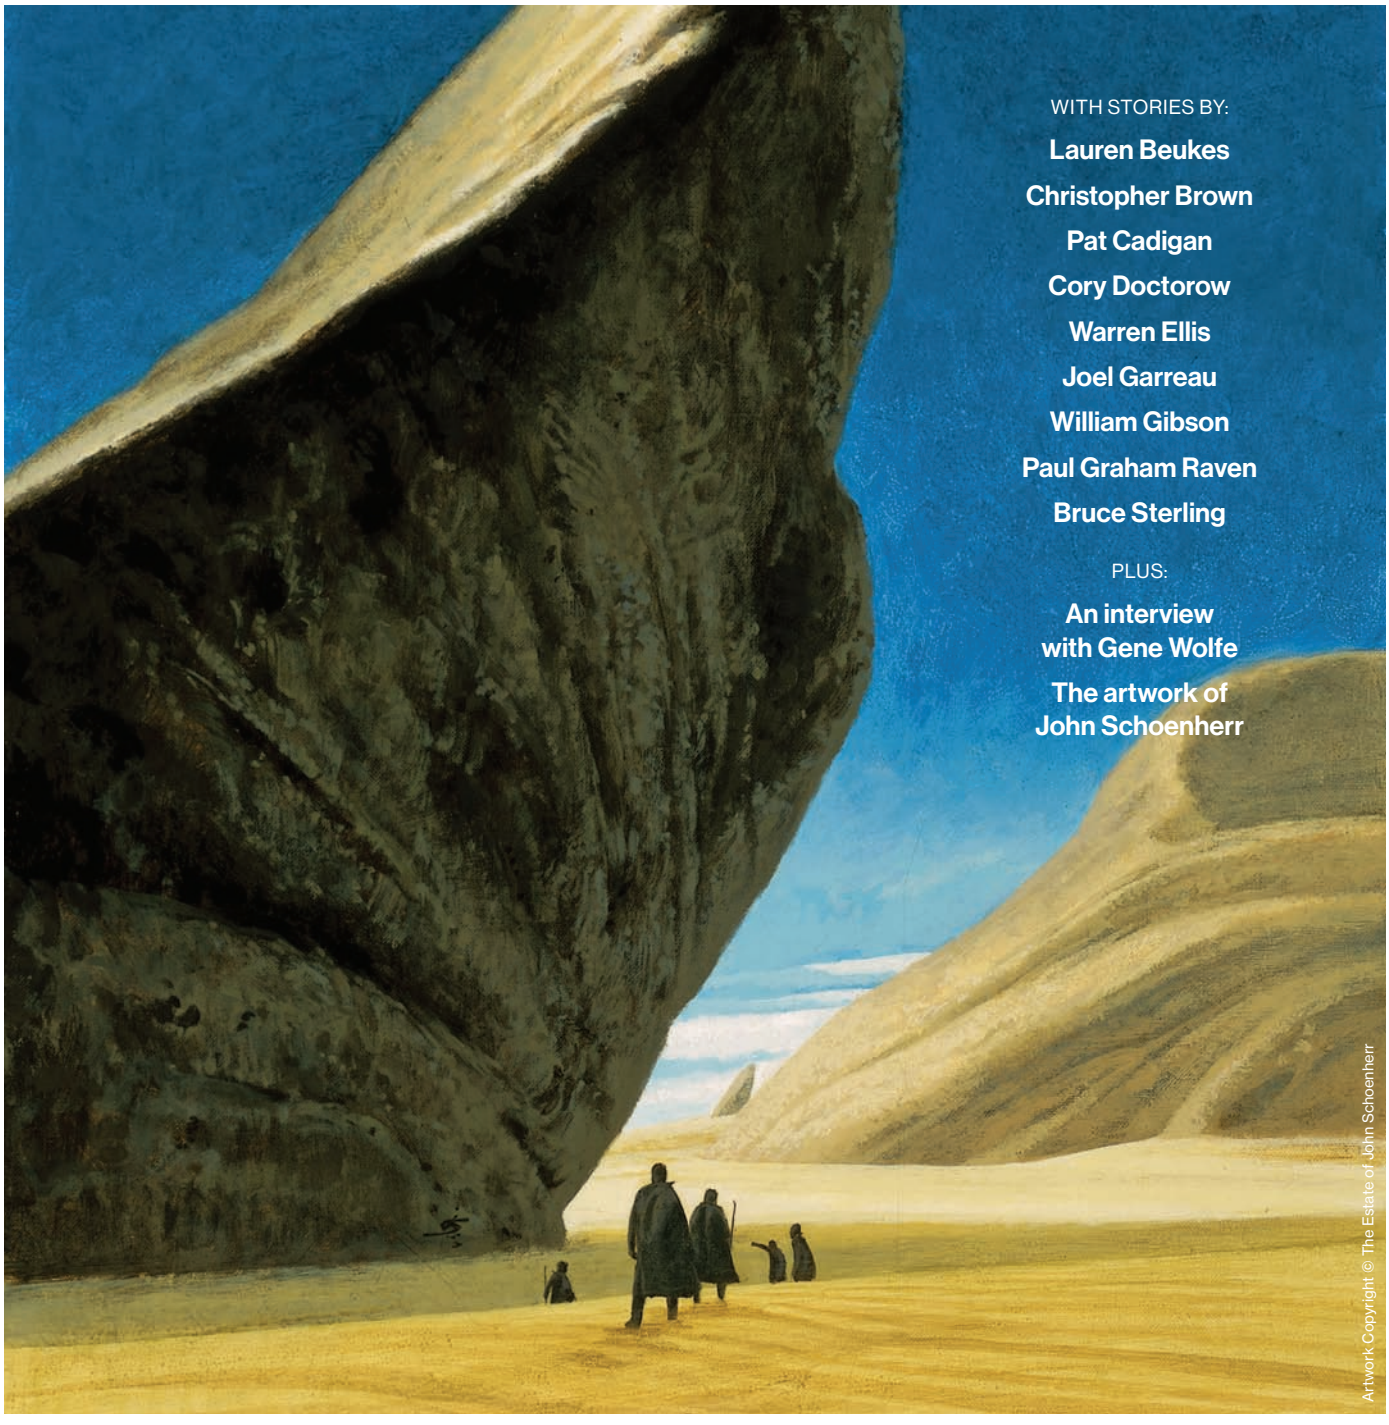

TWELVE TOMORROWS

The 2014 edition of MIT Technology Review's SF anthology

Visionary stories of the near future inspired by today's new technologies

Available for iPad, Kindle, and as a limited edition printed volume

Learn more: www.technologyreview.com/twelvetomorrows

MIT Technology Review

Upfront

2013, when it staged a kind of debutante moment at a leading research conference on machine learning. At Harrah's Casino on the shore of Lake Tahoe, DeepMind researchers showed off software that had learned to play three classic Atari games—Pong, Breakout, and Enduro—better than a human expert. The software was equipped only with access to the controls and display, a knowledge of the score, and an instinct to make that score as high as possible. It became an expert gamer entirely through trial and error.

DeepMind made use of a newly fashionable machine-learning technique called deep learning, which involves processing data through networks of crudely simulated neurons. But it combined deep learning with other tricks to make something with an unexpected level of intelligence.

"People were a bit shocked because they didn't expect that we would be able to do that at this stage of the technology," says Stuart Russell, a professor and artificial-intelligence specialist at the University of California, Berkeley.

In particular, DeepMind had used a technique called reinforcement learning, which is inspired by the work of animal psychologists such as B. F. Skinner. This led to software that learns by taking actions and receiving feedback on the results of those actions. Artificial-intelligence researchers have been tinkering with reinforcement learning for decades. But until DeepMind's Atari demo, no one had built a system capable of learning anything nearly as complex as how to play a computer game. One reason it was possible was a trick borrowed from Hassabis's favorite area of the brain. Part of the Atari-playing software's learning process involved replaying its past experiences over and over to try to extract the most accurate hints on what it should do in the future. "That's something that we know the brain does," says Hassabis.

Hassabis expects swift progress toward a very powerful new form of artificial intelligence.

"When you go to sleep, your hippocampus replays the memory of the day back to your cortex."

More than a year later, Russell and other researchers are still puzzling over exactly how that trick and others led to such remarkable results, and what else they might be used for. But Google didn't take long to recognize the importance of the technology, announcing a month after the Tahoe demonstration that it had acquired the company.

Today Hassabis leads what is now called Google DeepMind. The group was roughly 75 people strong at the time it joined Google, and around three-quarters of its members work on fundamental research. The rest form an "applied research team" that looks for opportunities to apply DeepMind's techniques to existing Google products. In addition, Hassabis dreams of creating "AI scientists" that could do things like generate and test new hypotheses about disease in the lab. When prodded, he says that DeepMind's software could also be useful to robotics, an area in which Google has recently invested heavily. "One reason we don't have more robots doing more helpful things is that they're usually preprogrammed," he says. "They're very bad at dealing with the unexpected or learning new things."

Hassabis's reluctance to talk about applications might be coyness, or it could be that his researchers are still in the early stages of understanding how to advance the company's AI software. One strong indicator that Hassabis expects swift progress toward a powerful new form of AI is that he is setting up an ethics board to consider the possible downsides of advanced artificial intelligence. "It's something that we or other people at Google need to be cognizant of. We're still playing Atari games currently," he says, laughing. "But we are on the first rungs of the ladder."

TO MARKET

On Button

Curie

COMPANY:
Intel

PRICE:
Not disclosed

AVAILABILITY:
Second half of 2015

Intel has developed a minuscule computer, called Curie, to power compact wearable devices. About the size of a button, Curie includes a Bluetooth radio, a six-axis motion sensor with accelerometer and gyroscope, and components that can rapidly and precisely differentiate between different types of physical activity. The device runs on a coin-size battery (although Intel hasn't said for how long). Intel's announcement is a sure sign that hardware makers are eager to build wearable devices of all kinds. It also points to the unwelcome size and bulk of many existing smart watches and smart glasses, a factor that has turned many would-be users away.

COURTESY INTEL

BECOME THE NEXT GREAT DISRUPTOR.

REALIZE YOUR IDEAS. CHANGE THE WORLD.

AUGMENTED REALITY

BIOTECH

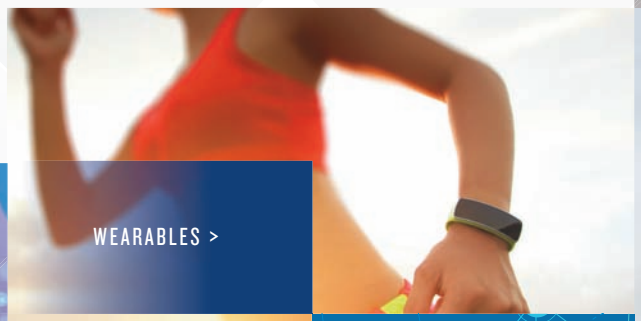

WEARABLES >

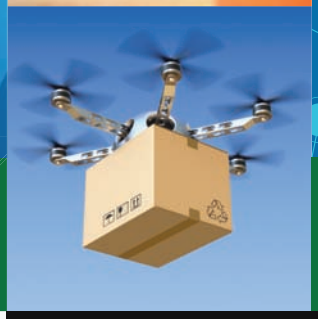

< DRONES

< THE INTERNET OF THINGS

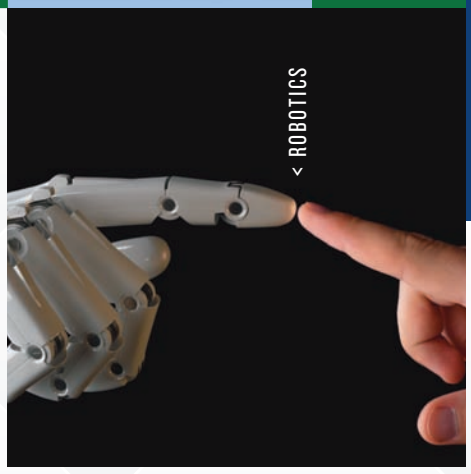

< ROBOTICS

FACTORY OF THE FUTURE

JABIL

WWW.JABIL.COM

Upfront

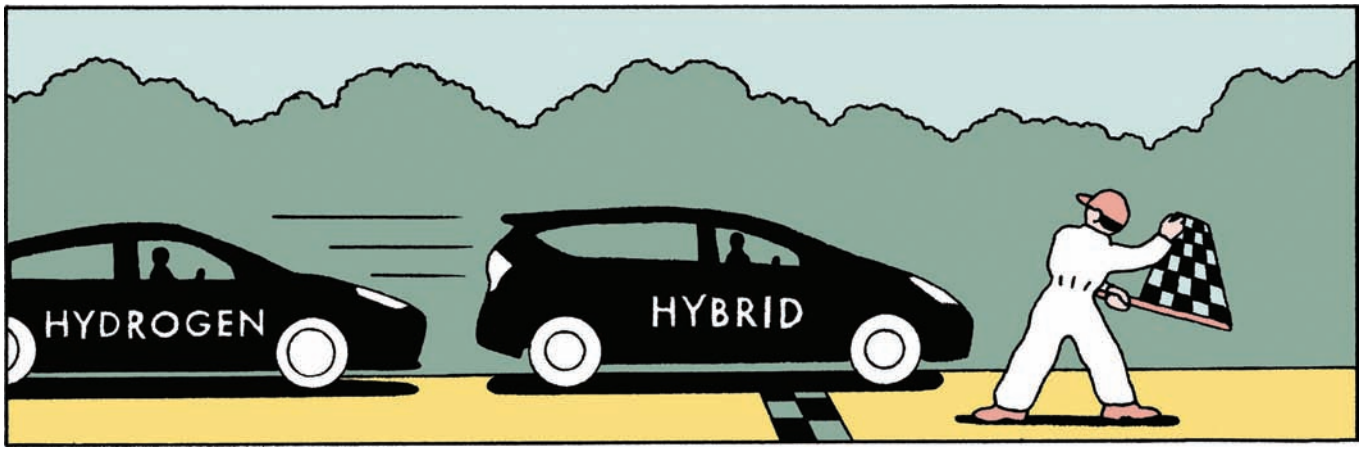

Forget Hydrogen Cars and Buy a Hybrid

Hybrids are a much more cost-effective way to reduce carbon emissions than new cars powered by fuel cells.

By Kevin Bullis

If you want to help cut greenhouse-gas emissions, you should probably skip the hydrogen-fuel-cell cars now coming to market and buy a (much cheaper) hybrid instead.

After decades of research and small-scale demonstrations, hydrogen cars are finally rolling into view. These vehicles use electric motors, but their electricity comes not from a battery but from hydrogen, processed in a chemical reaction that takes place inside a fuel cell.

Researchers and engineers have greatly lowered the costs of fuel cells—by as much as 95 percent—in recent years. That, along with pressure to meet emissions regulations, especially in California, means the technology is finally coming to market.

Last year, Hyundai started leasing hydrogen-powered Tucson Fuel Cell SUVs in California. Toyota launched a hydrogen car called the Mirai in Japan in December and said it would sell the car in the United States this year. Meanwhile, GM, Honda, and others are developing their own hydrogen vehicles.

Carmakers are keen to extol the environmental credentials of these new models. Hyundai advertises that its cars emit no carbon dioxide, while Toyota boasts that its hydrogen cars "leave nothing behind but water."

But these ads are a little misleading. Though the only thing that comes out of the cars' tailpipes is, indeed, water vapor, the hydrogen they run on is mostly made from natural gas in a process that releases significant amounts of greenhouse gases into the atmosphere.

An analysis by the Union of Concerned Scientists suggests that producing hydrogen from natural gas for the Hyundai Tucson Fuel Cell vehicle emits about as much carbon dioxide as a car that gets 38 miles per gallon. That's far better than the gasoline-powered version of the Tucson, which gets 25 miles per gallon. But you can buy a number of cars that get better than 38 miles per gallon.

Emerging technologies for producing hydrogen could eventually make fuel-cell cars cleaner and cheaper, though. Hydrogen can be made using renewable sources of electricity to power an electrolyzer, which splits water into its constituent hydrogen and oxygen atoms. The problem is that this is still far more costly than making hydrogen from natural gas.

It may also be possible to engineer catalysts to absorb sunlight and use its energy to split water. That would make hydrogen generation simpler and cheaper.

For now, though, the main advantage of hydrogen cars over electric cars is that they can be recharged more quickly. Even the fastest chargers available—for the Tesla Model S—take about 20 minutes to add 130 miles of range. You can fill a Hyundai's hydrogen tank, which holds enough for 265 miles, in 10 minutes.

On the other hand, while there are plans to install 40 public hydrogen-fueling stations next year, right now there are still only three in the entire United States.

ALAIN PILON

$18 billion

Apple's profit in its first fiscal quarter of 2015, more than Microsoft, IBM, Intel, Facebook, Yahoo, eBay, and Netflix combined.

A Small Step Toward Manufacturing Artificial Cells

A microfluidic cell copies some basic functions of life.

By Courtney Humphries

In a step toward sophisticated artificial cells, scientists have engineered a silicon chip that can produce proteins from DNA, the most basic function of life.

The system, though relatively simple, suggests a path to mimicking life with partly manufactured components, says Roy Bar-Ziv, a materials scientist at the Weizmann Institute of Science in Israel, who is leading the work.

Cells constantly create proteins from instructions coded in DNA sequences. The amount of each protein that's made is controlled by other genes, often in complicated feedback loops. Bar-Ziv calls his cell-on-a-chip "a new system allowing us to examine how genes are turned on and off outside the living cell."

The chips were created using a technique Bar-Ziv's lab developed several years ago to anchor DNA to silicon by first coating the surface with a light-activated chemical. The system, described in a *Science* paper in August by Bar-Ziv, his students Eyal Karzbrun and Alexandra Tayar, and Vincent Noireaux of the University of Minnesota, allowed the researchers to create a simple network of interacting genes.

Scientists can easily synthesize proteins from DNA in a test tube, but those reactions eventually fizzle out as proteins accumulate and synthesis slows. That has made it hard to create functioning genetic circuits—interacting networks of genes and proteins—outside cells. Bar-Ziv says his chip overcomes that problem by flushing away waste products. Also, by changing the lengths of the channels leading to each DNA compartment, he was able to control how fast the proteins manufactured in it diffused to other areas of the chip, influencing other reactions. "If you want to reconstitute the dynamic nature of genes going up and down, you have to have a mechanism for degrading what you make," he says.

Bar-Ziv's chip may eventually lead to applications in diagnostics, environmental sensing, or drug screening. Scientists say it could be used to test new genetic constructs before they're put into actual cells like bacteria.

The next step, says Bar-Ziv, is to create more complex patterns and larger networks. He hopes to eventually be able to control hundreds of different genes in thousands of artificial cells at once, allowing them to communicate with and influence one another, not unlike cells in a living organism. That's still a ways off, he admits, but "going from one transistor to billions didn't happen in a day."

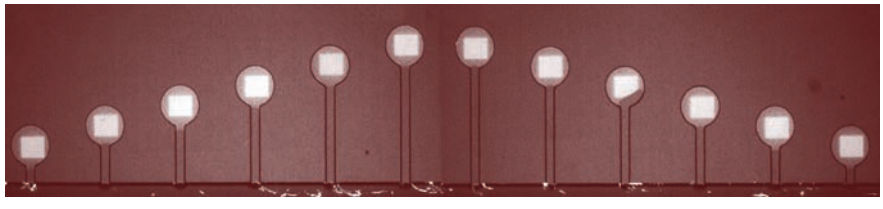

IMAGE COURTESY OF ALEXANDRA TAYAR AND EYAL KARZBRUN IN THE LAB OF PROFESSOR ROY BAR-ZIV

Chambers on the chip containing DNA are connected to cellular enzymes by microfluidic channels.

QUOTED

"We don't know what we are putting in. We don't really know what exactly it does, biologically."

—Freddie Fu, chairman of sports medicine at the University of Pittsburgh Medical Center, on orthopedic surgeons who give stem-cell treatments to professional athletes

"A model from the beginning of computing has been reflected in everything since, and it is holding us back."

—Kirk Bresniker, chief architect for HP's new computer, named The Machine

"We understand that they work, just not how they work. They can learn to do things that we can't even learn to do ourselves."

—Jeff Clune, an assistant professor of computer science at the University of Wyoming, speaking about software developed using a machine-learning technique known as deep learning

"Rocket made it ... but landed hard. Close, but no cigar this time."

—SpaceX CEO Elon Musk, in a tweet about his company's failed attempt to land the first stage of a Falcon 9 rocket on a floating barge after launch

Upfront

Pipe Dreams

Most people in the world are still not online (see "Project Loon," page 40), and many of those who are can't reap the full benefits of the global network.

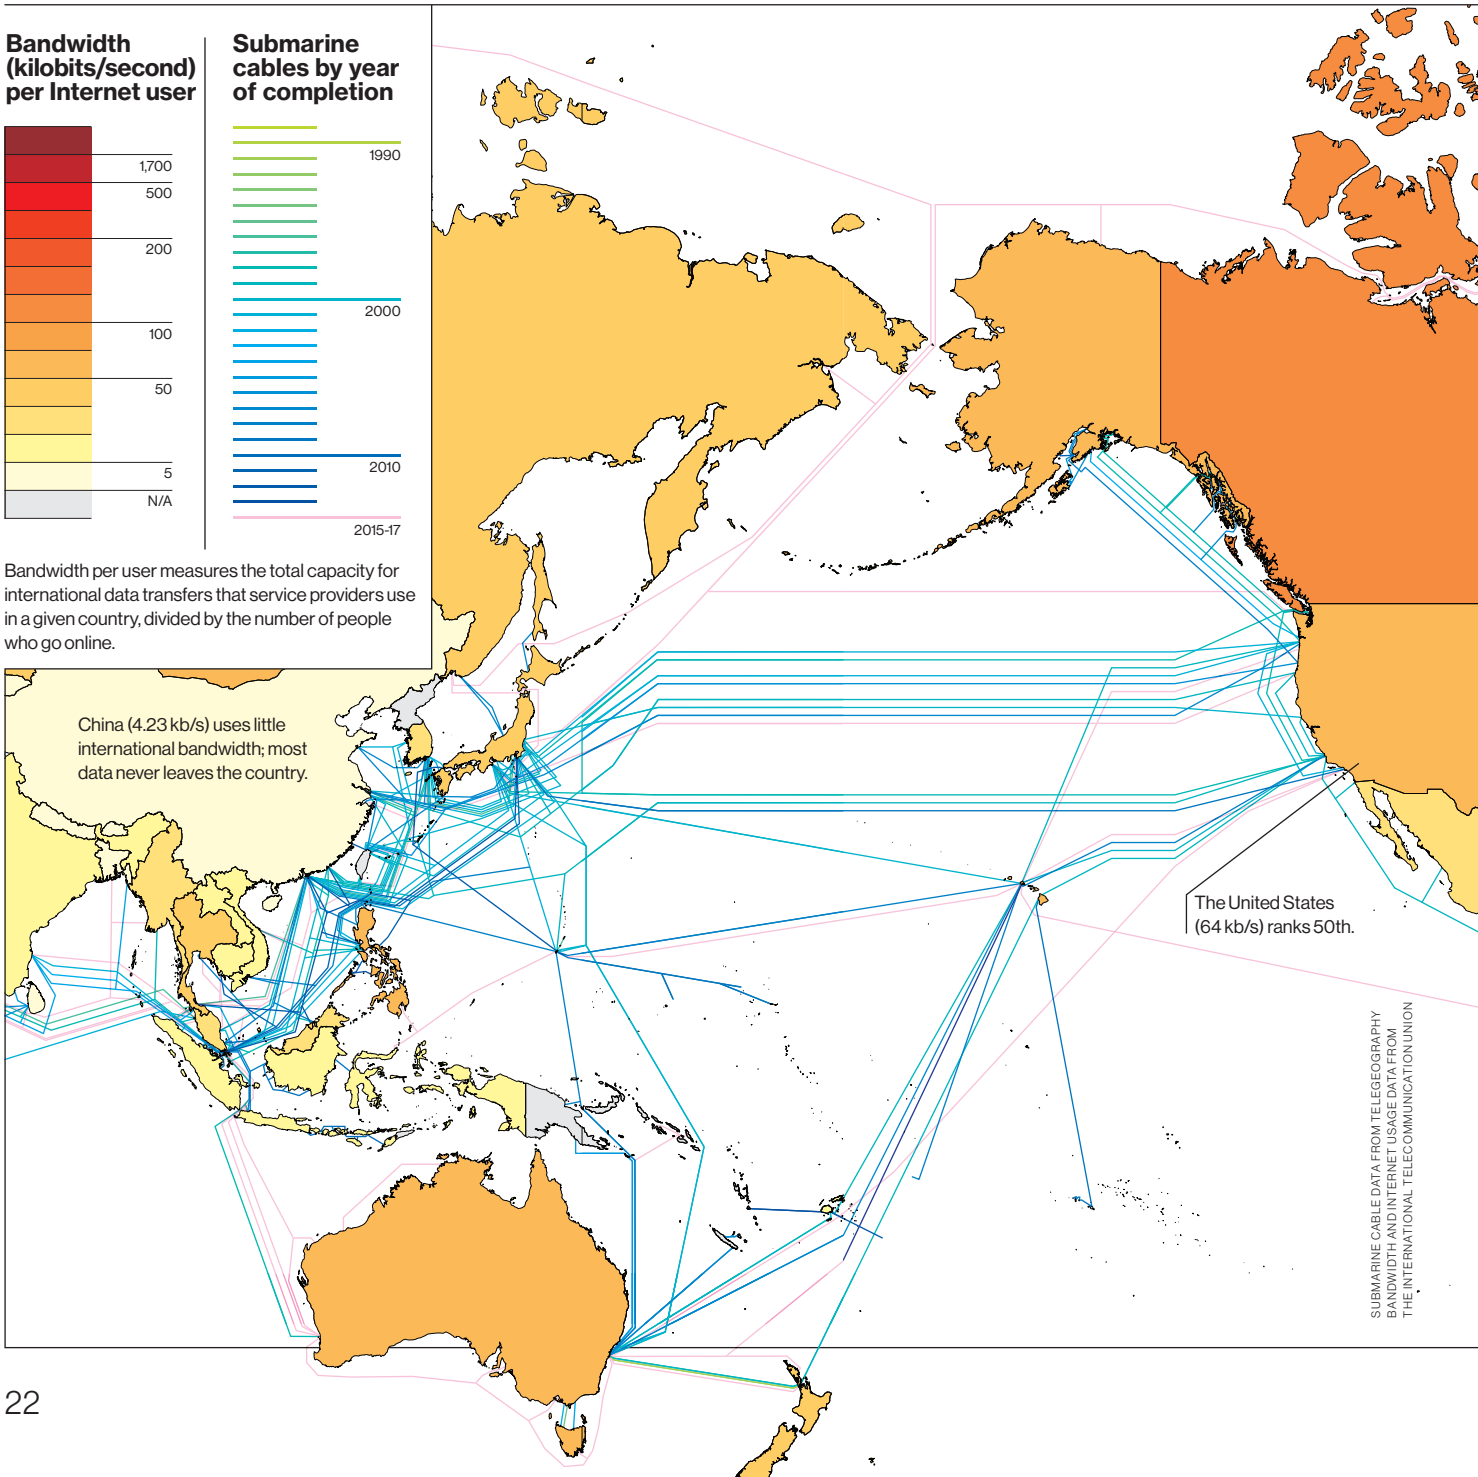

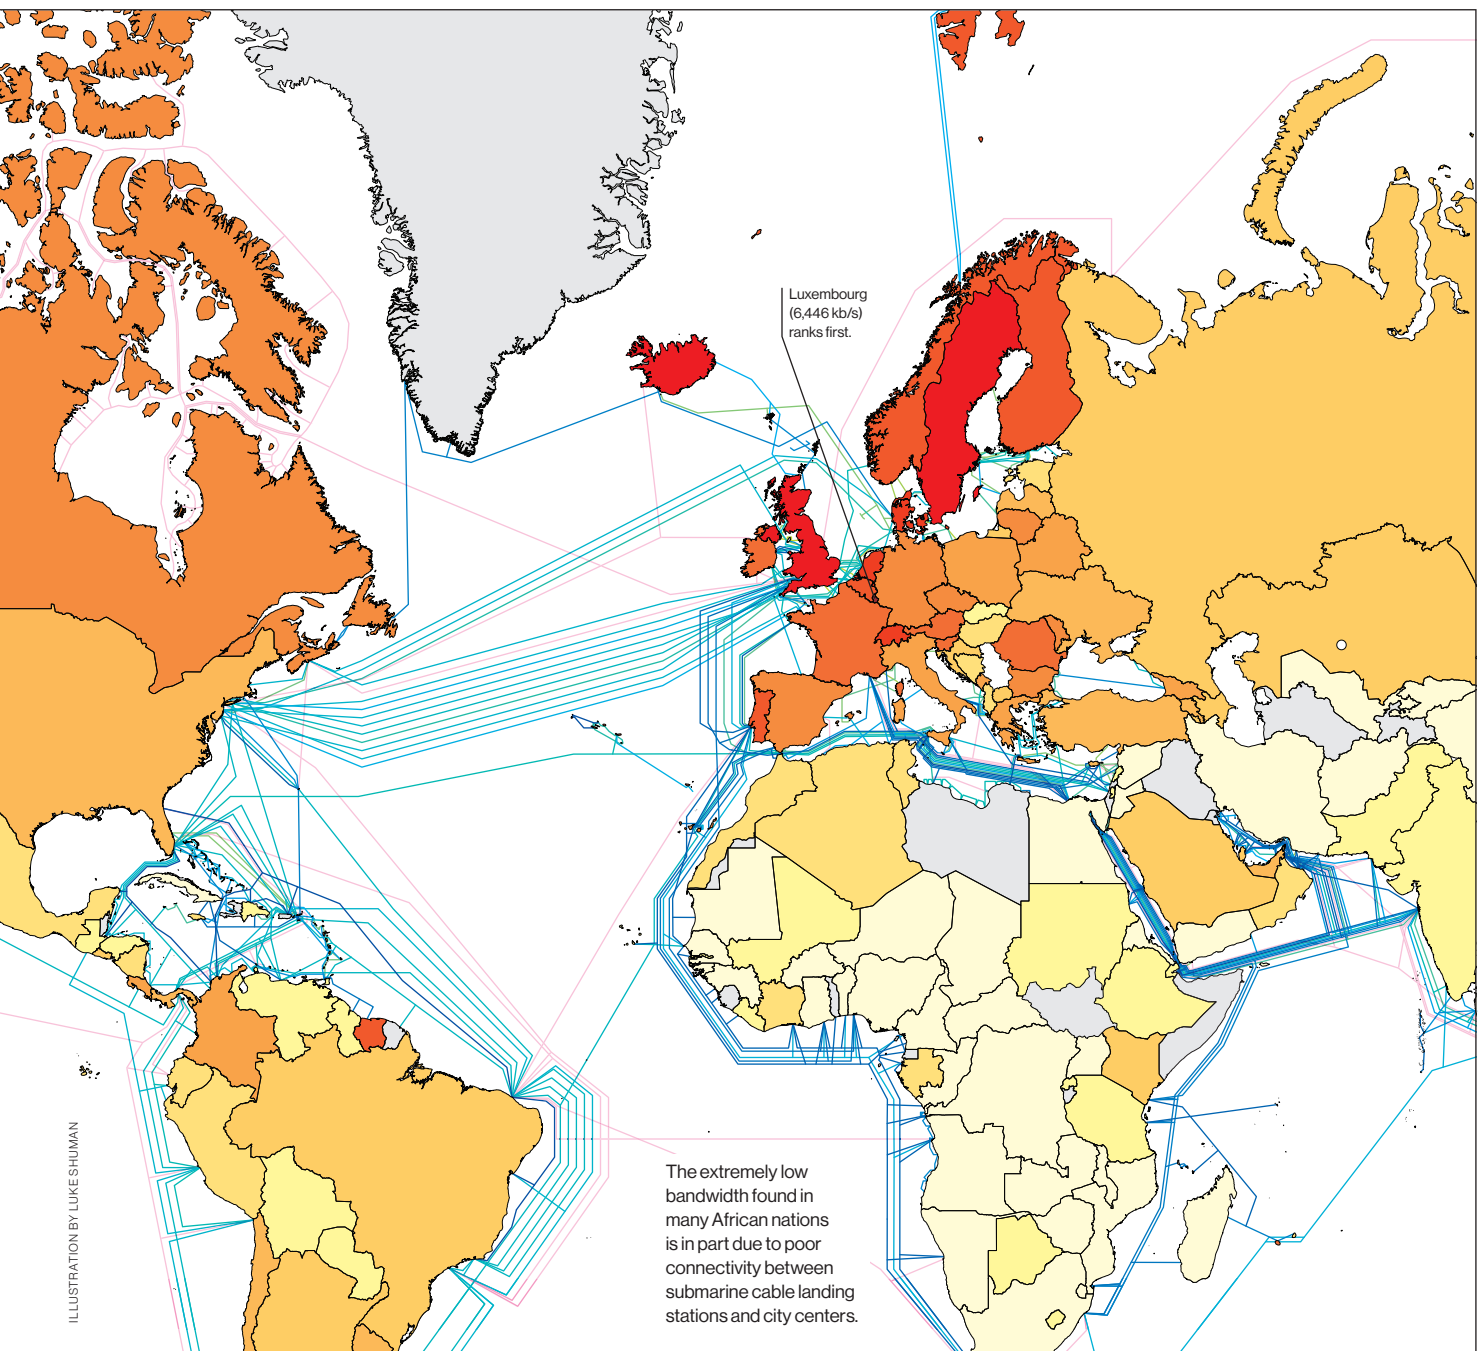

ILLUSTRATION BY LUKE SHUMAN

The extremely low bandwidth found in many African nations is in part due to poor connectivity between submarine cable landing stations and city centers.

Q+A

Steven Chu

As a leading and active scientist, Steven Chu broke the mold when he became energy secretary of the United States in 2009. In his four years of service, he made the Department of Energy more innovative, launching the Advanced Research Projects Agency for Energy to support projects not yet ready for private investment. He also created Innovation Hubs to bring people from different disciplines together on energy problems and rejuvenated funding for solar research. Chu, who shared the 1997 Nobel Prize in physics and directed the Lawrence Berkeley National Laboratory before his government appointment, is now restarting his research group at Stanford. In a conversation with David Talbot, chief correspondent of *MIT Technology Review*, he reflected on his time with the federal government and talked about the research and technology questions absorbing him today.

What left you the most frustrated or disappointed at the Department of Energy?
The press was sometimes frustrating to deal with. Often, reporters or their editors wanted to "make news" by generating controversy.

Inside the department, there was inertia to keep old programs unchanged, and friction against new approaches. For example, in research in biofuels, I wanted to cast a wide net for new ideas, but I was getting resistance against doing new research that didn't fit existing definitions of fuels listed by the Department of Agriculture. I wanted new ideas to be funded on the merits, and worry about categorization later.

What do you see as your biggest success and your biggest mistake?
My biggest success is that I was able to help recruit very capable scientists and engineers. Also, as a practicing scientist—during late nights or weekends—I was in a better position to ask the right questions. Perhaps my biggest mistake was to defer too much to "experts" on nonscientific matters at the beginning of my tenure. This was especially true if the advice was coming from handwringers who were more worried about negative reactions than doing the right thing.

What should President Obama try to get done on energy in his final two years?
President Obama, through the EPA, is doing the right thing by pushing on mercury, particulate matter, and carbon dioxide standards for all power plants above a certain size. I would like to have him begin a dialogue on policies for countries that have a meaningful price on carbon or are working to be less carbon-intensive in each particular industry.

For example, the carbon emission from the production of a particular grade of steel varies greatly. We need to think of how to prevent manufacturing and extraction industries from constantly migrating to the lowest-cost, most polluting producer. China is working hard to reduce the carbon intensity of its industries and is likely to put a price on carbon. I believe China and the U.S. can be leaders in starting this dialogue.

What projects excite you now?
After I left DOE, many companies asked me to join their boards of directors. I chose very few, including Amprius [a Stanford startup working on lithium-ion batteries]. Professor Yi Cui [a 2004 member of *MIT Technology Review*'s Innovators Under 35] and I brainstormed about new approaches to lithium-metal-anode batteries. We've published a couple of papers on new approaches. It's long been known that a lithium-metal–sulfur-cathode battery can potentially have five times higher energy density. We also seek a durable battery that can charge 10 times faster. Of course, as in all research, we may or may not succeed, but I think we have a shot.

You're also on the board of a Canadian startup called Inventys. Why?
I'm trying to help with some of the more technical aspects of capturing carbon from a natural-gas power plant—but also a coal, steel, or cement plant. Currently conventional methods that use amines [chemicals that absorb and then release carbon dioxide at different temperatures] are too expensive. We're hoping to reduce capture costs to $15 a ton for carbon dioxide; current technologies, when scaled, would cost around $60. Getting to $15 would make carbon capture feasible in the U.S. and China.

What's the fundamental physics breakthrough you'd most like to see?
Breakthroughs, by definition, are *unanticipated* surprises that lead to great things.

IMKE LASS/REDUX

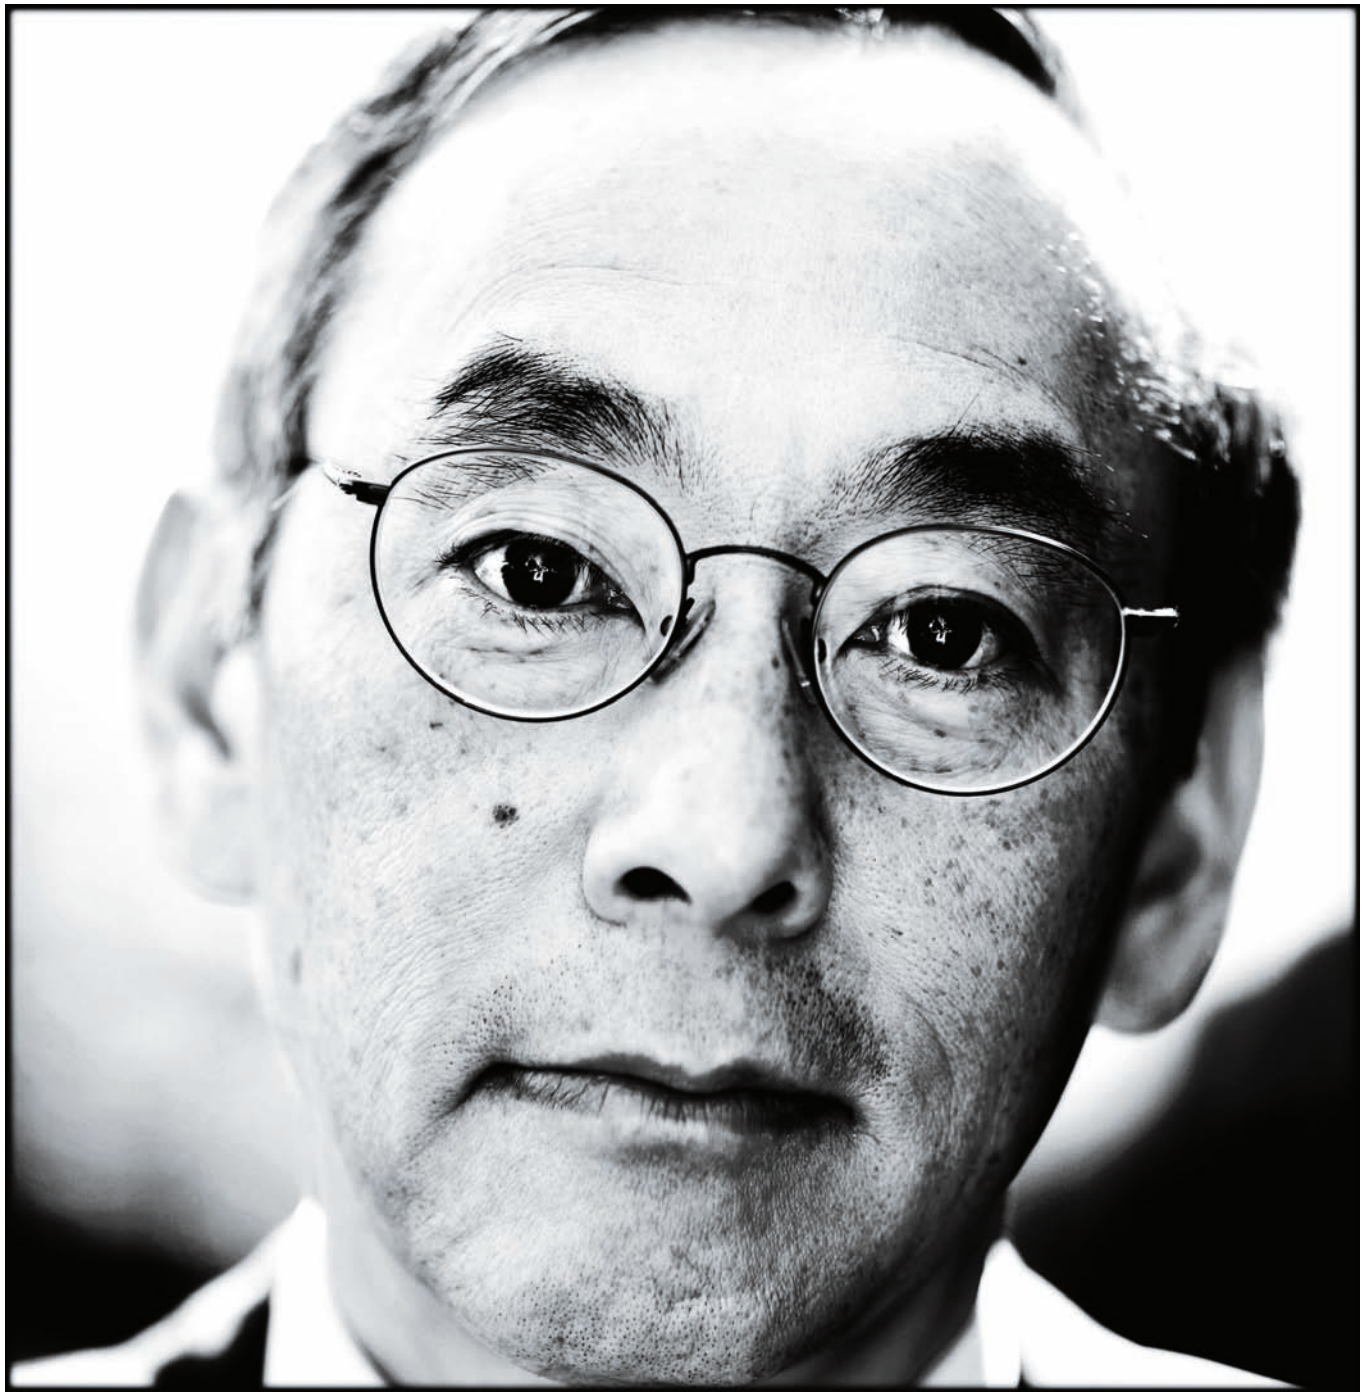

Who will industrialize the Internet of Things?

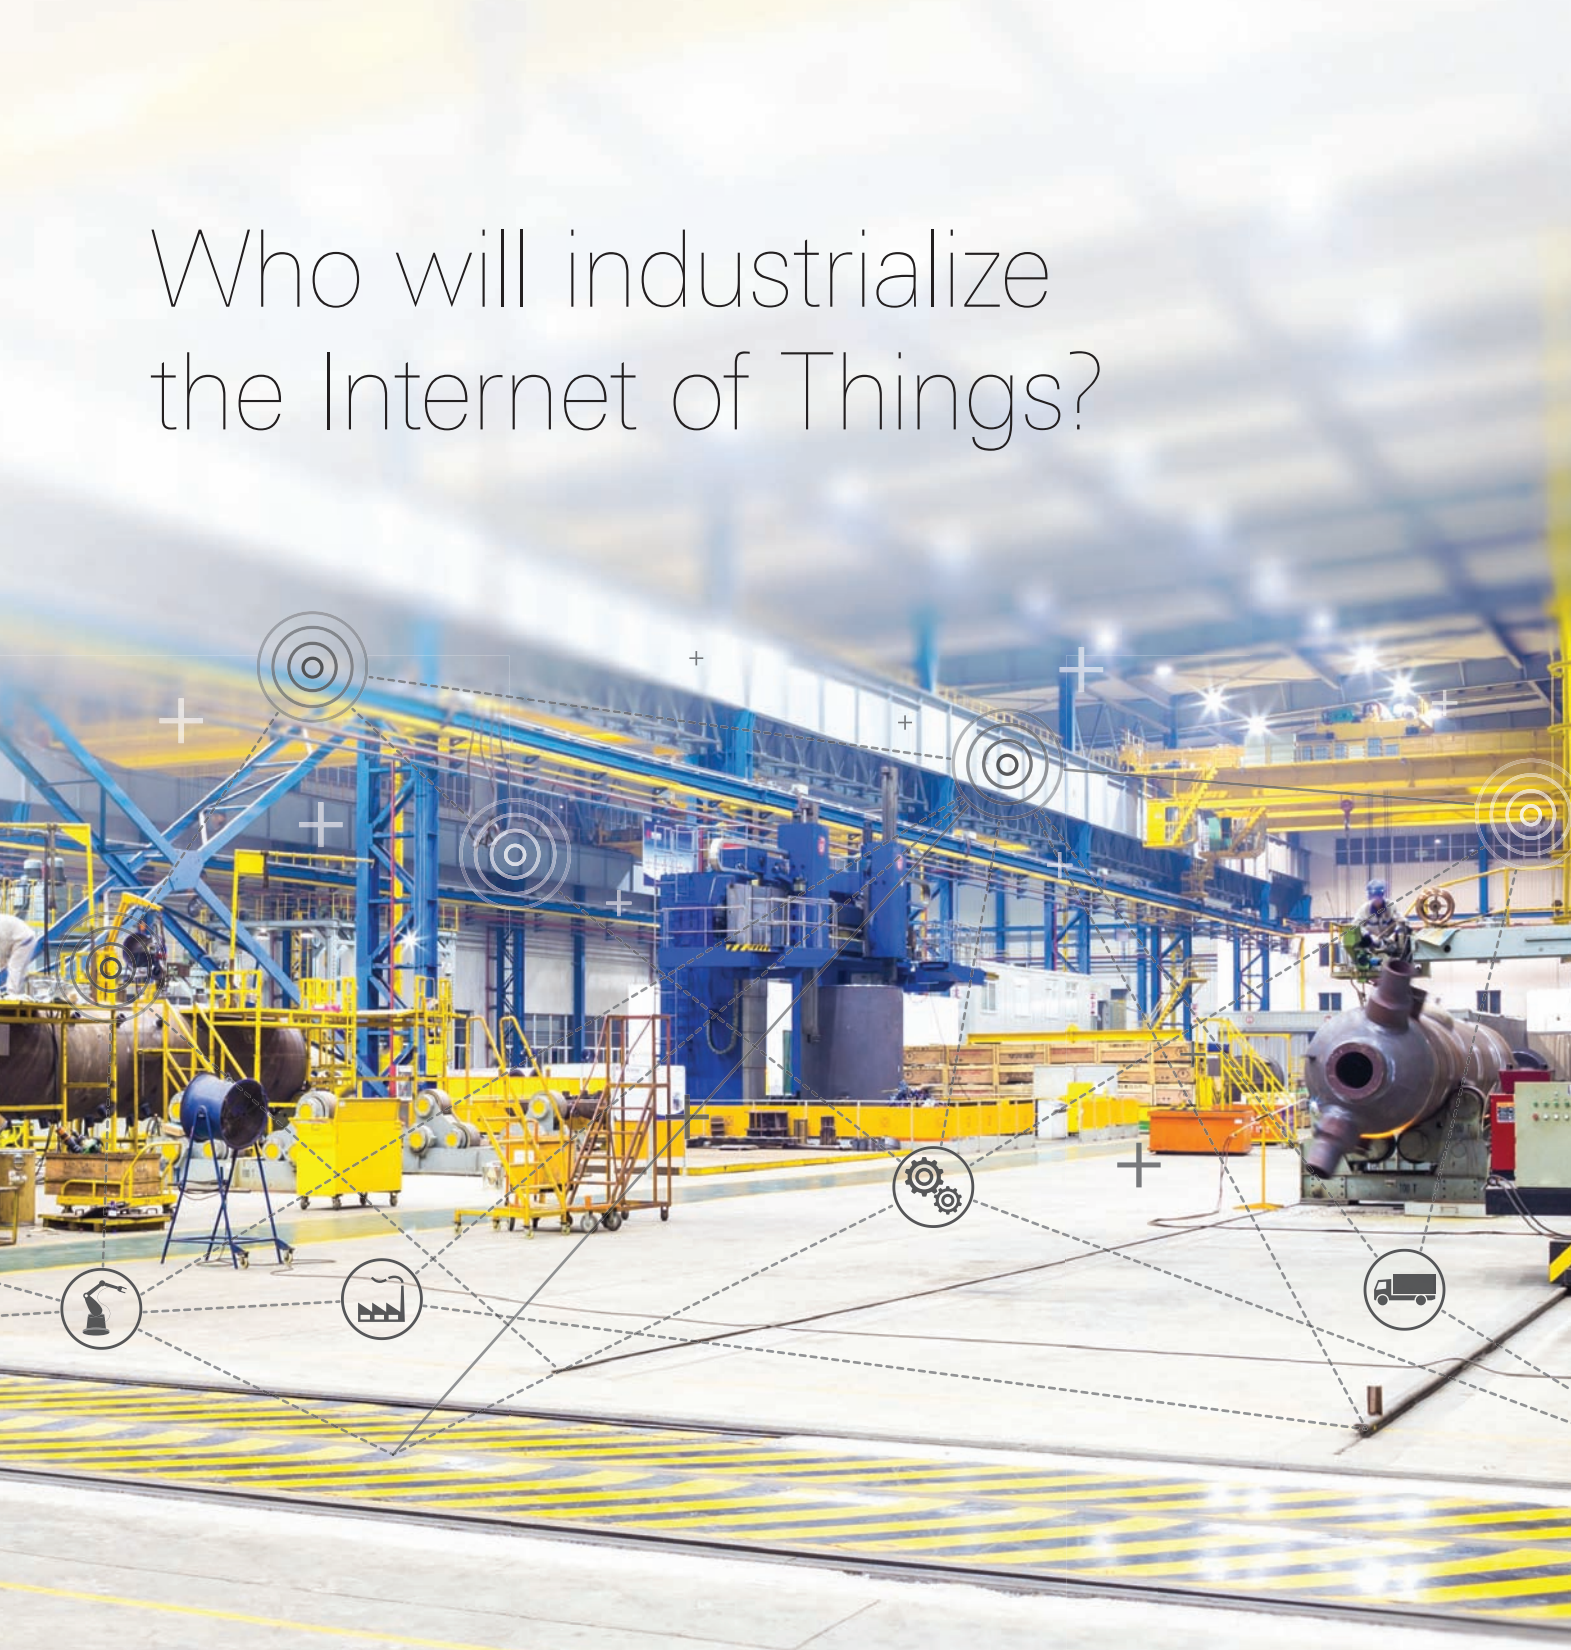

You and **NI** will. From smart manufacturing to the smart grid, bridging the digital and physical worlds is one of the most significant engineering opportunities we face today. NI, together with other visionary companies, is ensuring the superior design and performance of the increasingly complex systems that will fuel the Internet of Things and drive industrial adoption. Learn more at ni.com.

NATIONAL INSTRUMENTS™

ELLIOTT EARLS

10 Breakthrough Technologies

Not all breakthroughs are created equal. Some arrive more or less as usable things; others mainly set the stage for innovations that emerge later, and we have to estimate when that will be. But we'd bet that every one of the milestones on this list will be worth following in the coming years.

Magic Leap

A startup is betting more than half a billion dollars that it will dazzle you with its new approach to creating 3-D imagery.

By Rachel Metz

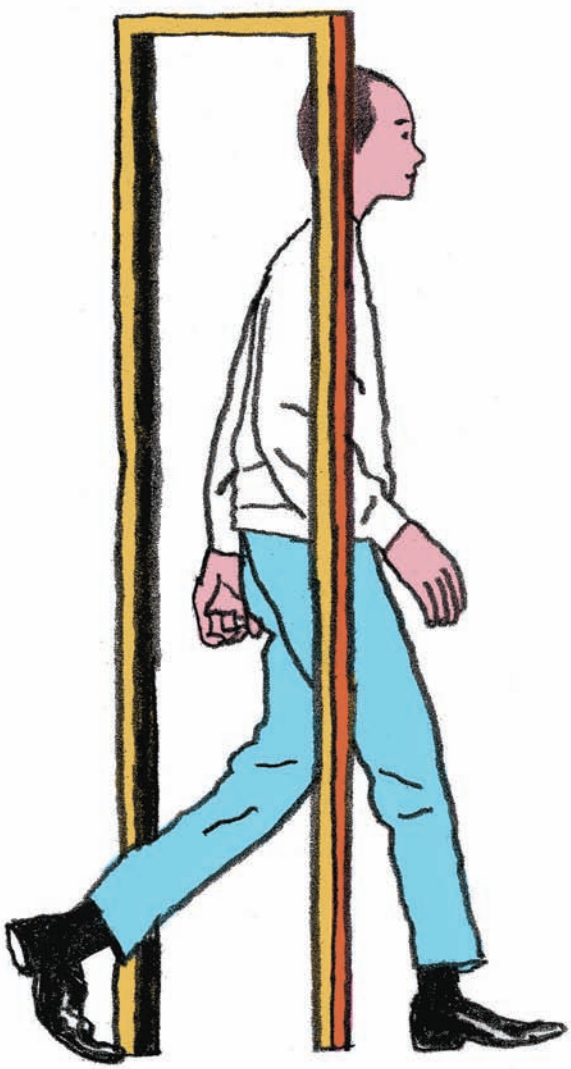

Logically, I know there isn't a hulking four-armed, twisty-horned blue monster clomping in circles in front of me, but it sure as hell looks like it.

I'm sitting behind a workbench in a white-walled room in Dania Beach, Florida, in the office of a secretive startup called Magic Leap. I'm staring wide-eyed through a pair of lenses attached to what looks like metal scaffolding that towers over my head and contains a bunch of electronics and lenses. It's an early prototype of the company's so-called cinematic-reality technology, which makes it possible for me to believe that the muscular beast with the gruff expression and two sets of swinging arms is actually in the room with me, hovering about seven feet in front of my face.

He's not just visible at a set distance. I'm holding a video-game controller that's connected to the demo station, and at the press of a button I can make the monster smaller or larger, move him right or left, bring him closer, or push him farther away.

Of course, I bring him as near as possible; I want to see how real he looks up close. Now he's about 30 inches from my

Breakthrough
A device that can make virtual objects appear in real life.

Why It Matters
The technology could open new opportunities for the film, gaming, travel, and telecommunications industries.

Key Players
- Magic Leap
- Microsoft

JOSH COCHRAN

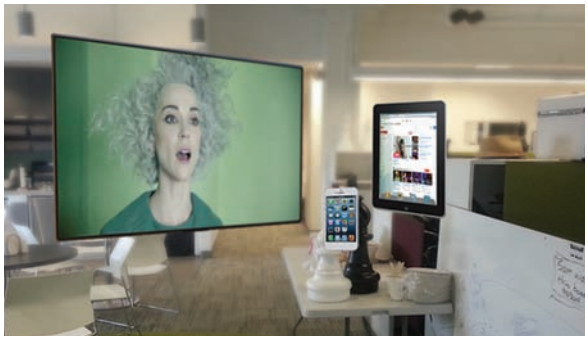

A video by the musician St. Vincent floats on a virtual screen in a break area in Magic Leap's headquarters.

eyeballs and, though I've made him pocket-sized, looks about as authentic as a monster could—he seems to have rough skin, muscular limbs, and deep-set beady eyes. I extend my hand to give him a base to walk on, and I swear I feel a tingling in my palm in expectation of his little feet pressing into it. When, a split second later, my brain remembers that this is just an impressively convincing 3-D image displayed in the real space in front of me, all I can do is grin.

Virtual- and augmented-reality technologies used in movies, smartphone apps, and gadgets tend to underdeliver on overhyped promises with images that look crappy. Typically that's because stereoscopic 3-D, the most commonly used method, is essentially tricking your eyes instead of working with the way you normally see things. It produces a sense of depth by showing each eye a separate image of the same object at a different angle. But since that forces you to look simultaneously at a flat screen in the distance and images that appear to be moving in front of you, it can make you dizzy and lead to headaches and nausea.

To be sure, stereoscopic 3-D has recently started getting better. The best system you can currently buy comes from Oculus VR, which Facebook purchased last spring for $2 billion; the $199 Gear VR, which was built in collaboration with Samsung and is aimed at software developers, lets you slide a Samsung smartphone into a headset to play games and watch videos.

But while Oculus wants to transport you to a virtual world for fun and games, Magic Leap wants to bring the fun and games to the world you're already in. And in order for its fantasy monsters to appear on your desk alongside real pencils, Magic Leap had to come up with an alternative to stereoscopic 3-D—something that doesn't disrupt the way you normally see things. Essentially, it has developed an itty-bitty projector that shines light into your eyes—light that blends in extremely well with the light you're receiving from the real world.

As I see crisply rendered images of monsters, robots, and cadaver heads in Magic Leap's offices, I can envision someday having a video chat with faraway family members who look as if they're actually sitting in my living room while, on their end, I appear to be sitting in theirs. Or walking around New York City with a virtual tour guide, the sides of buildings overlaid with images that reveal how the structures looked in the past. Or watching movies where the characters appear to be right in front of me, letting me follow them around as the plot unfolds. But no one really knows what Magic Leap might be best for. If the company can make its technology not only cool but comfortable and easy to use, people will surely dream up amazing applications.

Abovitz says he and his employees are trying to "blow away" their inner 11-year-olds.

COURTESY OF MAGIC LEAP

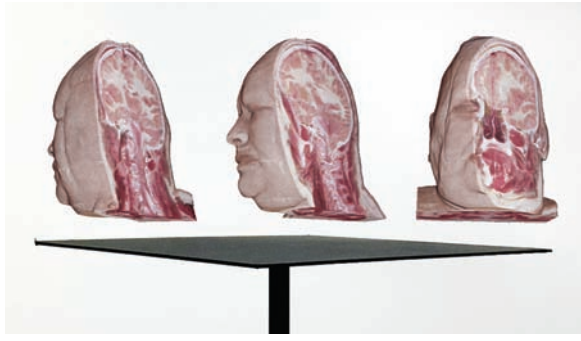

Top: In a demonstration of a medical or educational application, a cadaver head can be dissected one slice at a time. Bottom: A fake robot appears to stand on a real hand.

That's no doubt why Google took the lead in an astonishingly large $542 million investment round in Magic Leap last October. Whatever it is cooking up has a good chance of being one of the next big things in computing, and Google would be crazy to risk missing out. The investment looked especially prescient in January, when Microsoft revealed plans to release a sleek-looking headset this year. HoloLens, which lets you interact with holograms, sounds as if it's very similar to what Magic Leap is working on.

Behind the magic

Magic Leap won't say when it will release a product or how much the thing will cost, beyond that the price will be within the range of today's consumer mobile devices. When I press founder and CEO Rony Abovitz about such details, he'll only smile and say, "It's not far away."

He's sitting behind the desk in his office, which is just down the road from the Fort Lauderdale–Hollywood airport. The shelves are lined with toys and View-Masters—the plastic gadgets that let you look at pictures in 3-D. Abovitz, 44, is a bear of a guy with a kind smile, and when I meet him he's dressed in black Nikes, a long-sleeved shirt, and slacks, his graying curly hair topped with a yarmulke. He's thoughtful and composed, which I find somewhat surprising given that the only time I had seen him before was in a video of his talk at a TEDx event in 2012 in Sarasota, Florida. It featured two people dressed as furry creatures called "Shaggles," Abovitz walking on stage dressed as an astronaut, and unintelligible rock music. Though the talk, called "The Synthesis of Imagination," came off as performance art (perhaps even a mockery of a TED talk), he swears there is a coherent message embedded in it; figure it out, he says, and he'll give you a yo-yo.

By day, Abovitz is a technology entrepreneur with a background in biomedical engineering. He previously founded Mako Surgical, a company in Fort Lauderdale that makes a robotic arm equipped with haptic technology, which imparts a sense of touch so that orthopedic surgeons have the sensation of actually working on bones as they trigger the robot's actions. Mako was sold to a medical technology company, Stryker, for nearly $1.7 billion in 2013. By night, Abovitz likes to rock out. He sings and plays guitar and bass in a pop-rock band called Sparkydog & Friends. And as he tells it, Magic Leap has its origins in both the robotic-surgery company and his life as a musician.

Combining virtual reality with the physical world appealed to Abovitz even at Mako. Although the robotic-arm technology could give surgeons the sensation of touching their instruments to bones, Abovitz also wanted to let them see virtual

bones as they went about this work. Over and over, he says, he tried out head-mounted displays made by different companies, but he was disappointed with them all. "They were all just complete crap," he says. "You'd put it on and it would give you a headache and it was awful, and I was wondering, 'Why is this so bad?'"

At the same time, Abovitz also wanted to take Sparkydog & Friends on a virtual tour. In U2's 1987 video for "Where the Streets Have No Name," the group, in a nod to an earlier move by the Beatles, plays an impromptu show on the roof of a Los Angeles liquor store. Abovitz yearned for his band to be able to do that, but virtually, and on a thousand rooftops at once.

About four years ago, he started mulling the problem over with John Graham Macnamara, a high school friend who had dropped out of Caltech's theoretical physics program. They became captivated by the idea of displaying moving holograms like the one in *Star Wars*. Holograms—3-D images that can be viewed from many angles—are made by accurately re-creating light fields, the patterns made when light rays bounce off an object. But Abovitz figured it would cost a lot and take lots of time to project even low-resolution holographic images. At one point, he remembers muttering, "There is no display that can actually work."

The next morning, though, he awoke with an idea: why bother with the painstaking steps needed to send a hologram out into a room for multiple people to see at once? Why not, instead, essentially make a hologram that only *you* see, doing it in a way that is natural for the eyes and brain to perceive, unlike stereoscopic 3-D? "We're spending half a billion dollars–plus to effectively make nothing happen to you, physiologically," Abovitz says.

The solution he and Macnamara and the rest of Magic Leap's team have come up with is still largely under wraps, and on the record they avoid discussing how the technology works except in vague terms, citing concerns about competition. But it's safe to say Magic Leap has a tiny projector that shines light onto a transparent lens, which deflects the light onto the retina. That pattern of light blends in so well with the light you're receiving from the real world that to your visual cortex, artificial objects are nearly indistinguishable from actual objects.

If the company can get this to work in a head-mounted display, showing images near the eyes and consistently refocusing them to keep everything looking sharp, it will make 3-D images much more comfortable to view, says Gordon Wetzstein, an assistant professor of electrical engineering at Stanford who researches computational imaging and displays. "If they do what people suspect they do," Wetzstein says, "it will be amazing."

The Forerunners—and the Competition

1838: Sir Charles Wheatstone invents the first stereoscope, which uses two angled mirrors to reflect a separate image into each eye. It gives viewers the sense that they're seeing one image in three dimensions.

1922: A silent 3-D film, *The Power of Love*, is released; viewers wear glasses with two different-colored lenses—red and green—to watch it.

1961: Philco employees build the first known head-mounted display, called Headsight, which features a helmet outfitted with a cathode-ray tube and magnetic head-position tracking.

1962: Morton Heilig receives a patent for the Sensorama, a big, boxy machine that shows short 3-D films on a small, one-person display, combined with sensations like smell and wind to make the experience immersive.

1985: Jaron Lanier, who is credited with coining the term "virtual reality," founds VPL Research. It sells such products as the Data Glove, which lets people use their hands to interact with a virtual environment, and the EyePhone, a head-mounted display.

1990: Boeing scientists Thomas Caudell and David Mizell build a wearable, see-through display that can superimpose lines on a board—aimed at making it easier for workers to assemble bundles of wires on boards that will then be installed on an airplane.

2010: Quest Visual releases Word Lens, an application that makes it possible to point a smartphone camera at a sign written in Spanish and have it appear in English on the screen.

2012: Palmer Luckey raises $2.4 million on the crowdfunding site Kickstarter for his stereoscopic 3-D virtual-reality gaming headset, Oculus Rift. Two years later, Facebook will buy Oculus for $2 billion.

2015: Months after Google invests in Magic Leap, Microsoft shows off the HoloLens, which also uses a technology other than stereoscopic 3-D to make virtual objects appear to be integrated with the real world. It plans to release the gadget later in the year.

From virtual to reality

Magic Leap is working feverishly to get to that point. Since building its first prototype in 2011, the company has continued to shrink its technology down.

Already it works on something smaller than the unwieldy scaffolding I used. In another demonstration, using hardware on a cart, I can poke at a tiny flying steampunk robot, a character from a first-person-shooter game called Dr. Grordbort's Invaders that Magic Leap is making with Weta Workshop, which created many of the special effects in the *Hobbit* movies. The robot can follow my finger around with surprising accuracy, right between the cubicles in Magic Leap's office.

To judge from a look I get at a design prototype—a realistic-looking piece of hardware that's completely nonfunctional—the company appears to be aiming to fit its technology into a chunky pair of sports sunglasses wired to a square pack that fits into your pocket. A somewhat similar image in a patent application Magic Leap filed in January suggests as much, too. The company won't say for sure, though; Abovitz confirms that the headset will be a glasses-like wearable device, but I have to twist his arm to get him to agree to use even that hazy phrasing on the record.

It's clear that getting the technology into that small form will be very hard. The smallest demo hardware I've seen at Magic Leap can't yet match the experience of the bigger demo units. It includes a projector, built into a black wire, that's smaller than a grain of rice and channels light toward a single see-through lens. Peering through the lens, I spy a crude green version of the same four-armed monster that earlier seemed to stomp around on my palm. In addition to improving the resolution of smaller units, Magic Leap will have to cram in sensors and software that will track your eyes and fingers, so you can control and interact with its virtual creatures—which themselves will have to incorporate real-life objects into whatever they appear to be doing.

That's where last year's half-billion dollars of investment come in. Magic Leap is hiring like crazy. It's looking for software engineers for everything from eye tracking and iris recognition to the branch of artificial intelligence known as deep learning. It needs optical engineers, game designers, and other people who will dream up virtual objects to display. To give you a sense of where their minds might go, I saw ray guns and magic wands lying around the office. As its chief futurist, Magic Leap has hired the science fiction author Neal Stephenson, whose 1992 novel *Snow Crash* imagined a virtual world called the Metaverse.

The excitement of such quick growth is palpable at Magic Leap's brightly decorated headquarters, where staid office trappings are punctuated by red high-backed love seats and yellow chairs. Employees energetically describe the games, sensors, and ray guns they're working on.

With the massive investment last year, interest in the company has intensified. Abovitz says, "We went from 'Does anyone care about this?' to 'Okay, people do care.'" Now he and the team are feeling the weight of these expectations. He says, "The inner 11-year-old—we want to blow that away."

Abovitz was enigmatic in his brief appearance on a TEDx stage in 2012. "A few awkward steps for me; a magic leap for mankind," he said from inside his spacesuit.

COURTESY OF TEDX

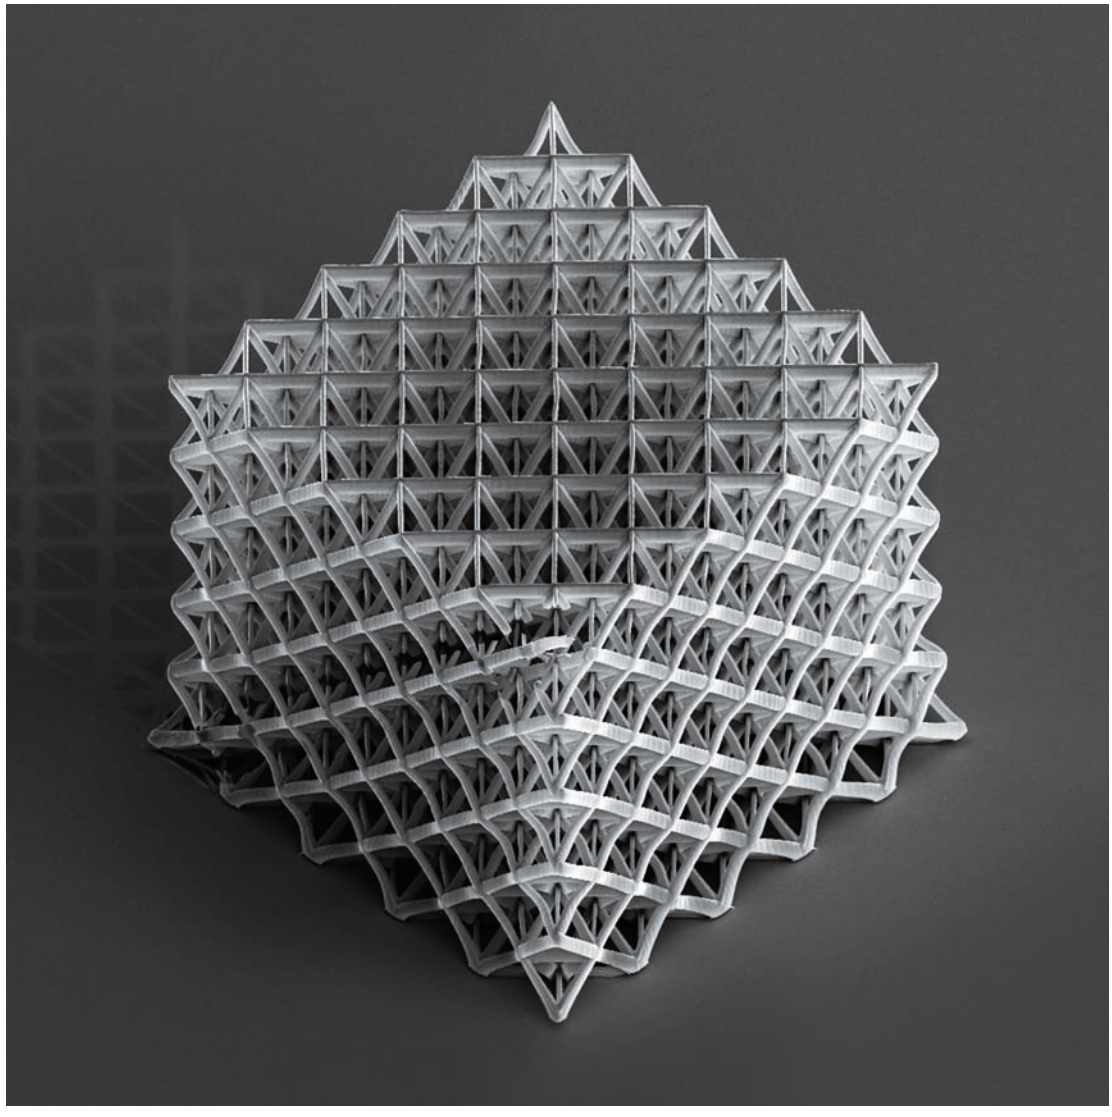

This ceramic cube, about 50 micrometers on each side, is light because it is mostly air.

COURTESY OF JULIA R. GREER

Nano-Architecture

A Caltech scientist creates tiny lattices with enormous potential.

By Katherine Bourzac

Breakthrough
Materials whose structures can be precisely tailored so they are strong yet flexible and extremely light.

Why It Matters
Lighter structural materials would be more energy-efficient and versatile.

Key Players
- Julia Greer, Caltech
- William Carter, HRL Laboratories
- Nicholas Fang, MIT
- Christopher Spadaccini, Lawrence Livermore National Laboratory

To visit the lab of Caltech materials scientist Julia Greer is to enter a realm where the ordinary rules of physical stuff don't seem to apply. Greer designs and builds nanomaterials that behave in ways surprising to those of us who spend our days in a world where strong materials like ceramic and steel tend to be heavy, while lightweight ones are weak. When Greer controls architecture at the nanoscale, the rules change.

Conventional ceramics are strong, heavy, and (as anyone who has dropped a plate knows) brittle, prone to shattering. But last year Greer created a ceramic that is one of the strongest and lightest substances ever made. It's also not brittle. In a video Greer made, a cube of the material shudders a bit as a lab apparatus presses down hard on it, then collapses. When the pressure is removed, it rises back up "like a wounded soldier," she says. "It's unreal, isn't it?" Greer often rushes to meetings around campus on Rollerblades and talks so fast that she demands focused listening. Peering into this beautiful, otherworldly nanolattice on her computer screen, she slows down for a while.

If materials like Greer's could be produced in large quantities, they could replace composites and other materials used in a wide range of applications, because they'd be just as strong at a fraction of the weight. Another possibility is to greatly increase the energy density of batteries—the amount of power they can hold at a given size. To do that, researchers have been trying to develop electrodes that are lighter than the ones used in today's batteries but can store more energy. However, promising electrode materials such as silicon are prone to cracking under strain. An electrode made by coating a metal nanolattice with silicon could have crack-resistant toughness in its very structure.

Greer in her Caltech lab, holding a model of the atomic structure of a metal.

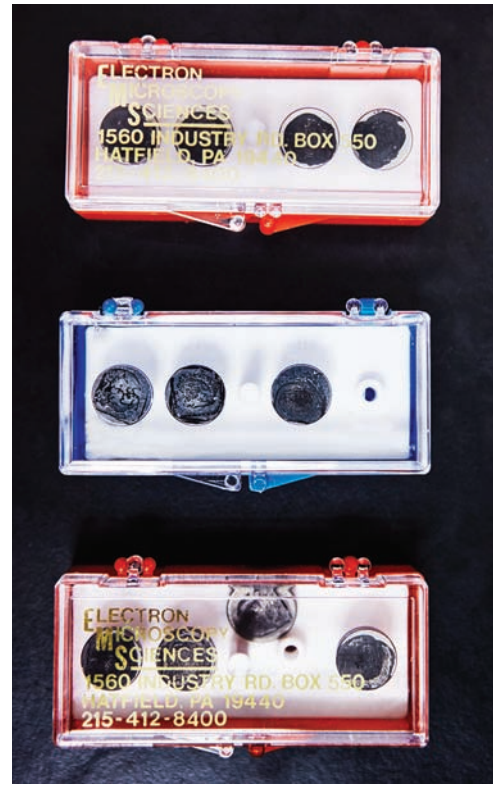

Top right: These disks are used to store nanoscale lattices. Bottom right: A scanning electron microscope has an arm that can compress and bend nanostructures.

Photographs by Anais & Dax

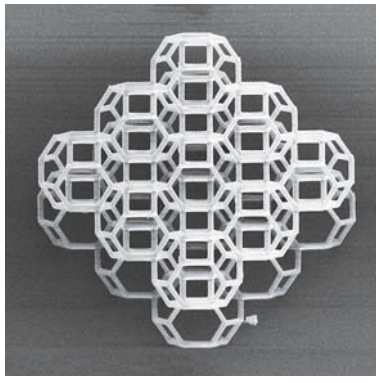

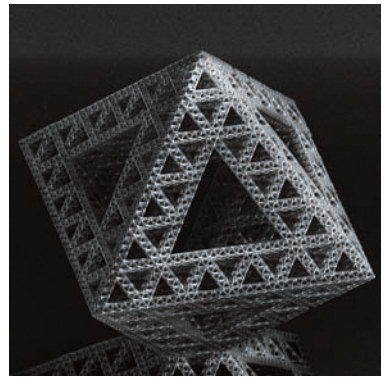

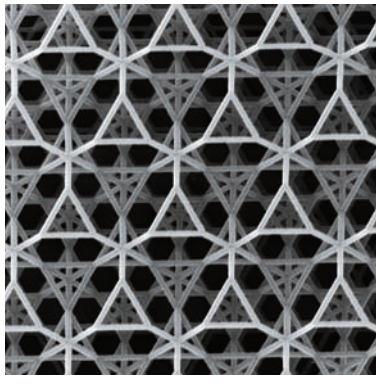

Fine-tuning materials' architecture at the nanoscale yields distinctive patterns—and unusual properties.

COURTESY OF JULIA R. GREER

The key to creating such wondrous materials is an arsenal of specialized machines—some of which Greer has rebuilt to suit her purposes—that make it possible to precisely control structure at the nanoscale over relatively large areas. Greer jogs down two floors of stairs to the basement lab where she keeps these precision instruments to isolate them from vibrations. One machine, found behind two heavy black curtains, is a sort of 3-D printer that uses flashes of laser light to very slowly build intricate polymer scaffolds. A student of Greer's coats the polymer with metals, ceramics, or other materials and then shaves off the sides, making it possible to etch away the polymer inside. The result is a little block of material made up of nanoscale trusses crisscrossed like the struts in the Eiffel Tower—but each strut's walls are only about 10 nanometers thick.

Without Greer's method, building something like this is impossible. She shows me a sample that came about from an earlier collaboration with researchers at HRL Laboratories in Malibu, California, who are producing materials with larger, microscale trusses. It's made out of nickel and looks somewhat like a metal scouring sponge. When she lets it drift onto my palm, I can barely feel it touch down, and the subversion of expectations is confusing. This metal is, literally, lighter than a feather. It could make for ultralight thermal insulation—an application her HRL colleagues are pursuing.

The featherweight nickel shows the promise of architectural control in making new materials with weird properties. But it's also a reminder of how far Greer has to go in scaling up her methods: so far, she can't make enough of the nanostructured materials to cover your palm.

Greer is determined to use her nanofabrication methods for a variety of materials, and a long list of collaborators are interested in their unusual properties. She can space the nanoscale walls in light-emitting materials or thermal insulation to precisely control the flow of light or heat. She's working with two battery makers to use her nanostructures to study electrochemistry. And she is teaming with biologists to see whether the nanostructured ceramic could serve as a scaffold for growing bones—such as the tiny ones in the ear whose degeneration is one cause of deafness.

In hopes of making such applications feasible, she is working to speed up the high-resolution laser-printing process. Greer has a six-millimeter-square fleck of the nanostructured ceramic she made last year. It is about as thick as a sheet of paper but took about a week to make.

"For us to do scientific experiments, we don't need a large amount," she says. "The question now is: how do you scale this?"

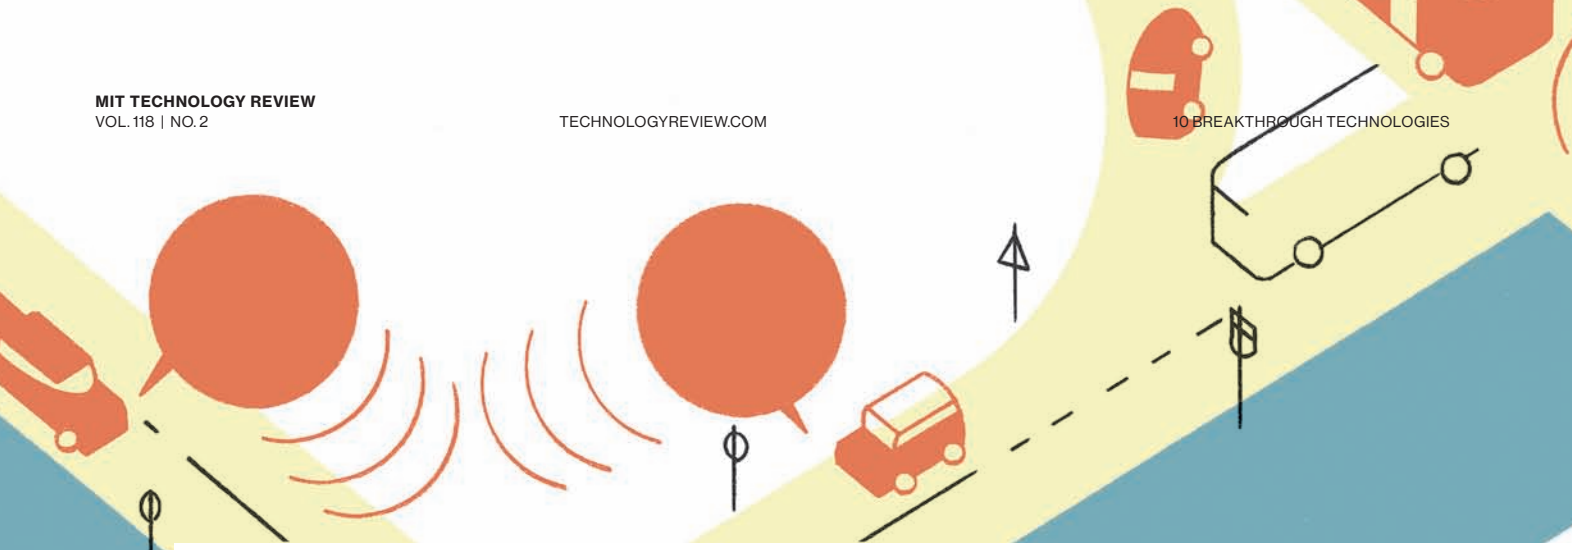

Car-to-Car Communication

A simple wireless technology promises to make driving much safer.

By Will Knight

Breakthrough
Cars that can talk to each other to avoid crashes.

Why It Matters
More than a million people are killed on roads worldwide every year.

Key Players
- General Motors
- University of Michigan
- National Highway Transportation Safety Administration

Hariharan Krishnan hardly looks like a street racer. With thin-rimmed glasses and a neat mustache, he reminds me of a math teacher. And yet on a sunny day last September, he was speeding, seemingly recklessly, around the parking lot at General Motors' research center in Warren, Michigan, in a Cadillac DTS.

I was in the passenger seat as Krishnan wheeled around a corner and hit the gas. A moment later a light flashed on the dashboard, there was a beeping sound, and our seats started buzzing furiously. Krishnan slammed on the brakes, and we lurched to a stop just as another car whizzed past from the left, its approach having been obscured by a large hedge. "You can see I was completely blinded," he said calmly.

The technology that warned of the impending collision will start appearing in cars in just a couple of years. Called car-to-car or vehicle-to-vehicle communication, it lets cars broadcast their position, speed, steering-wheel position, brake status, and other data to other vehicles within a few hundred meters. The other cars can use such information to build a detailed picture of what's unfolding around them, revealing trouble that even the most careful and alert driver, or the best sensor system, would miss or fail to anticipate.

Already many cars have instruments that use radar or ultrasound to detect obstacles or vehicles. But the range

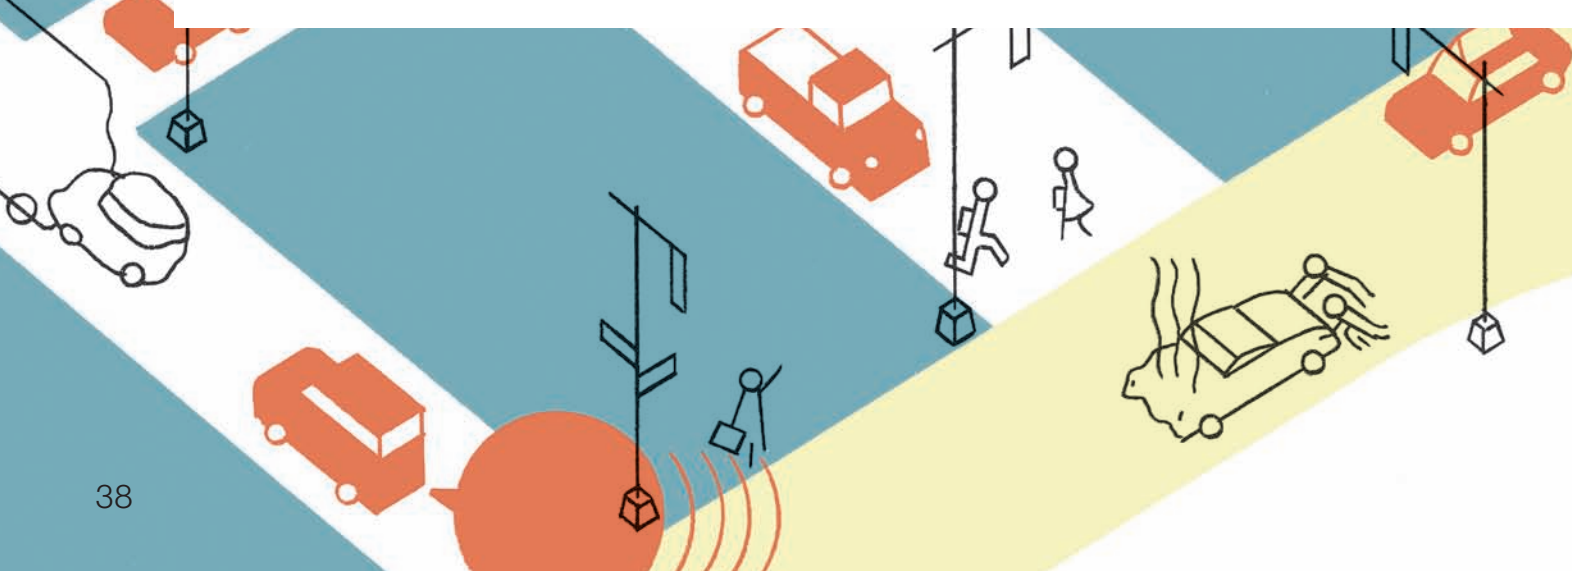

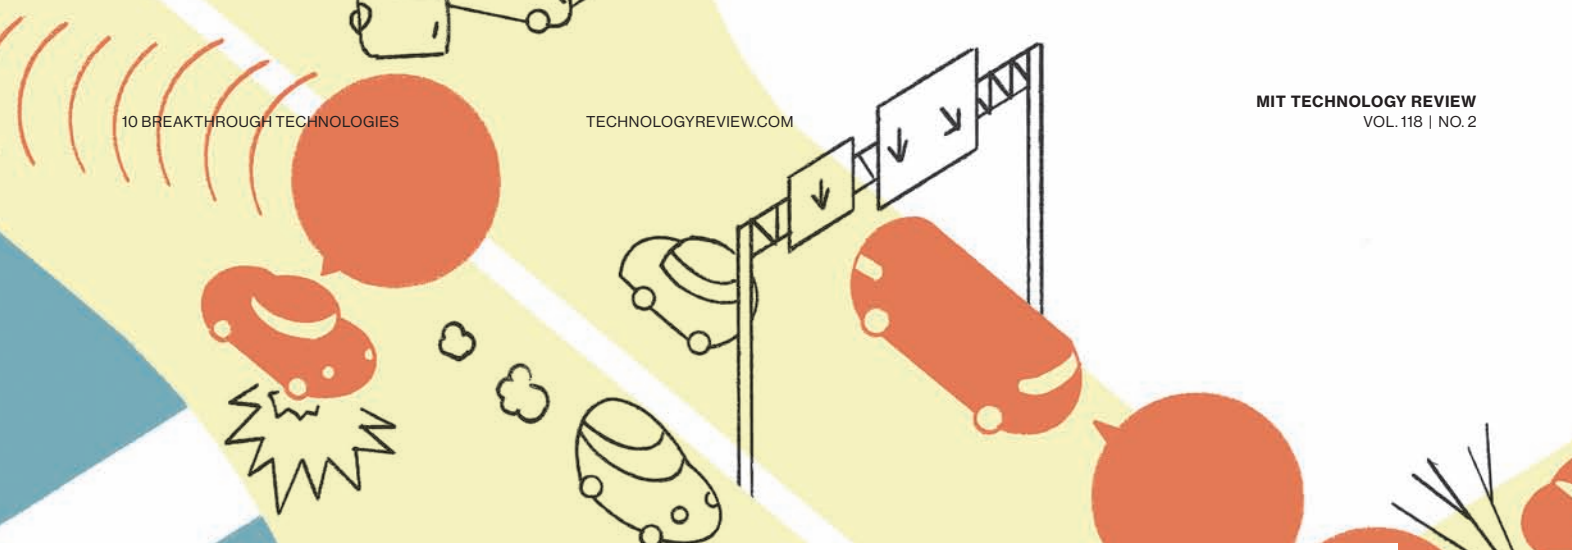

of these sensors is limited to a few car lengths, and they cannot see past the nearest obstruction.

Car-to-car communication should also have a bigger impact than the advanced vehicle automation technologies that have been more widely heralded. Though self-driving cars could eventually improve safety, they remain imperfect and unproven, with sensors and software too easily bamboozled by poor weather, unexpected obstacles or circumstances, or complex city driving. Simply networking cars together wirelessly is likely to have a far bigger and more immediate effect on road safety.

Creating a car-to-car network is still a complex challenge. The computers aboard each car process the various readings being broadcast by other vehicles 10 times every second, each time calculating the chance of an impending collision. Transmitters use a dedicated portion of wireless spectrum as well as a new wireless standard, 802.11p, to authenticate each message.

Krishnan took me through several other car-to-car safety scenarios in the company's parking lot. When he started slowly pulling into a parking spot occupied by another car, a simple alert sounded. When he attempted a risky overtaking maneuver, a warning light flashed and a voice announced: "Oncoming vehicle!"

More than five million crashes occur on U.S. roads alone every year, and more than 30,000 of those are fatal. The prospect of preventing many such accidents will provide significant impetus for networking technology.

Just an hour's drive west of Warren, the town of Ann Arbor, Michigan, has done much to show how valuable car-to-car communication could be. There, between 2012 and 2014, the National Highway Transportation Safety Administration and the University of Michigan equipped nearly 3,000 cars with experimental transmitters. After studying communication records for those vehicles, NHTSA researchers concluded that the technology could prevent more than half a million accidents and more than a thousand fatalities in the United States every year. The technology stands to revolutionize the way we drive, says John Maddox, a program director at the University of Michigan's Transportation Research Institute.

Shortly after the Ann Arbor trial ended, the U.S. Department of Transportation announced that it would start drafting rules that could eventually mandate the use of car-to-car communication in new cars. The technology is also being tested in Europe and Japan.

There will, of course, also be a few obstacles to navigate. GM has committed to using car-to-car communication in a 2017-model Cadillac. Those first Cadillacs will have few cars to talk to, and that will limit the value of the technology. It could still be more than a decade before vehicles that talk to each other are commonplace.

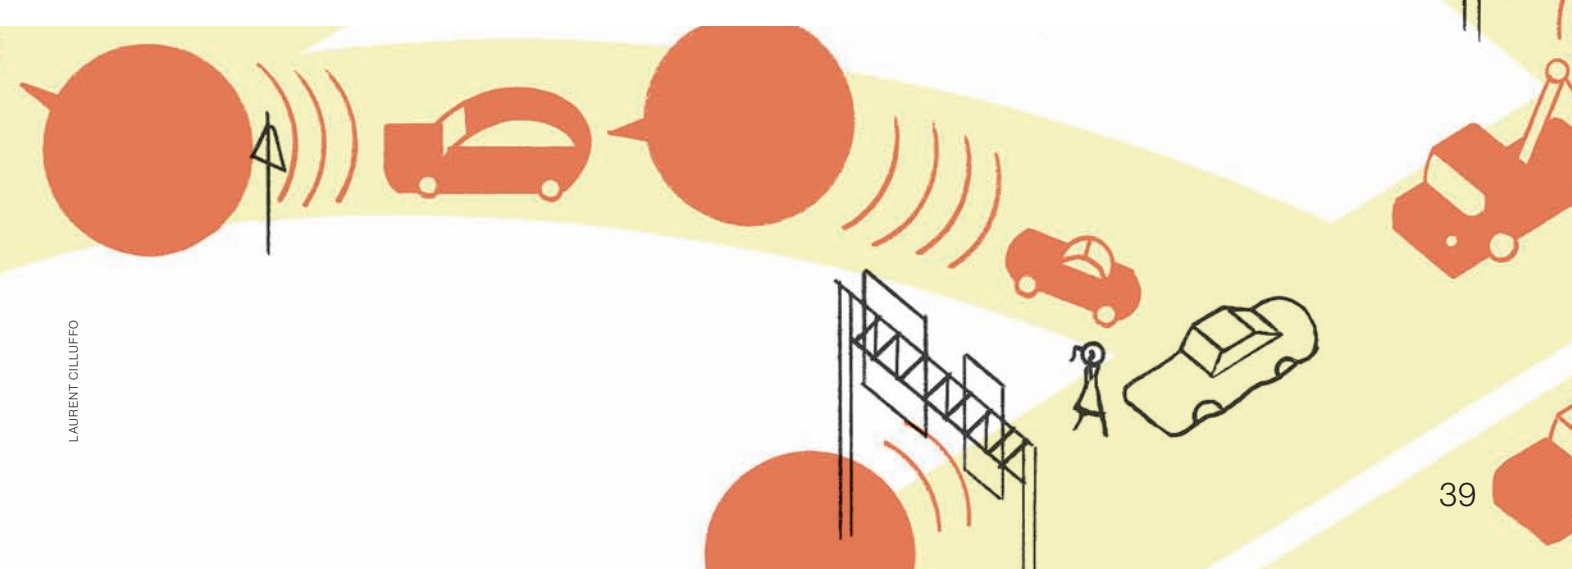

LAURENT CILLUFFO

Project Loon

Billions of people could get online for the first time thanks to helium balloons that Google will soon send over many places cell towers don't reach.

By Tom Simonite

Photographs by RC Rivera

MIT TECHNOLOGY REVIEW
VOL. 118 | NO. 2

You climb 170 steps up a series of dusty wooden ladders to reach the top of Hangar Two at Moffett Federal Airfield near Mountain View, California. The vast, dimly lit shed was built in 1942 to house airships during a war that saw the U.S. grow into a technological superpower. A perch high in the rafters is the best way to appreciate the strangeness of something in the works at Google—a part of the latest incarnation of American technical dominance.

On the floor far below are Google employees who look tiny as they tend to a pair of balloons, 15 meters across, that resemble giant white pumpkins. Google has launched hundreds of these balloons into the sky, lofted by helium. At this moment, a couple of dozen float over the Southern Hemisphere at an altitude of around 20 kilometers, in the rarely visited stratosphere—nearly twice the height of commercial airplanes. Each balloon supports a boxy gondola stuffed with solar-powered electronics. They make a radio link to a telecommunications network on the ground and beam down high-speed cellular Internet coverage to smartphones and other devices. It's known as Project Loon, a name chosen for its association with both flight and insanity.

Google says these balloons can deliver widespread economic and social benefits by bringing Internet access to the 60 percent of the world's people who don't have it. Many of those 4.3 billion people live in rural places where telecommunications companies haven't found it worthwhile to build cell towers or other infrastructure. After working for three years and flying balloons for more than three million kilometers, Google says Loon balloons are almost ready to step in.

It is odd for a large public company to build out infrastructure aimed at helping the world's poorest people. But in addition to Google's professed desires to help the world, the economics of ad-supported Web businesses give the company other reasons to think big. It's hard to find new customers

Breakthrough
A reliable and cost-effective way to beam Internet service from the sky to places lacking it.

Why It Matters
Internet access could expand educational and economic opportunities for the 4.3 billion people who are offline.

Key Players
- Google
- Facebook

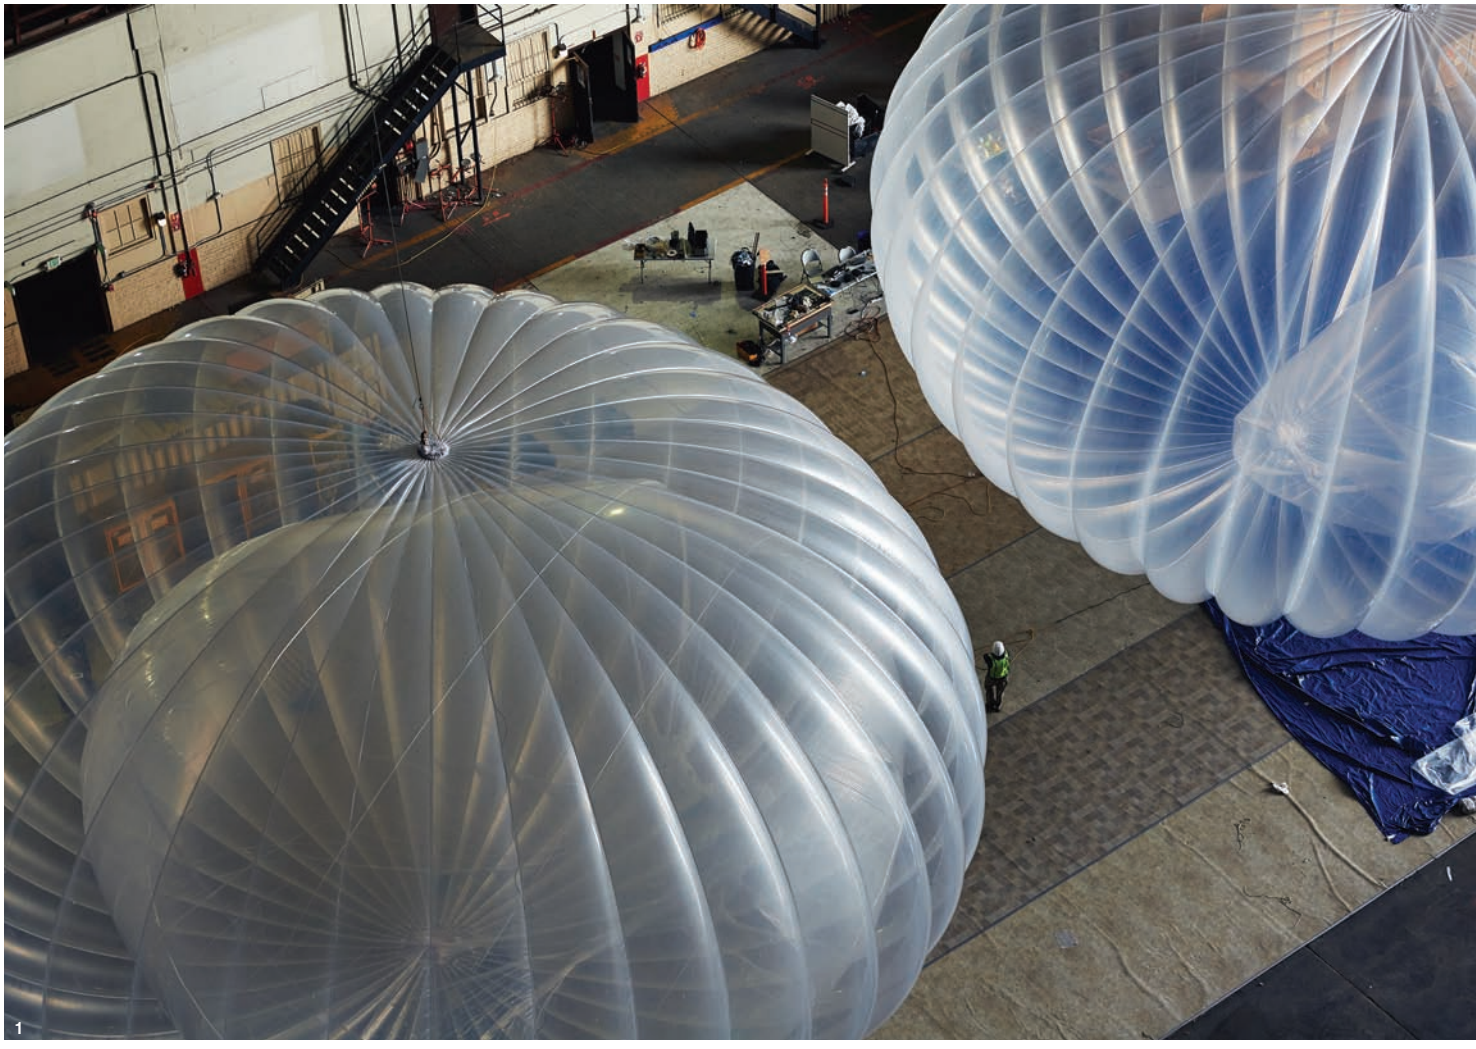

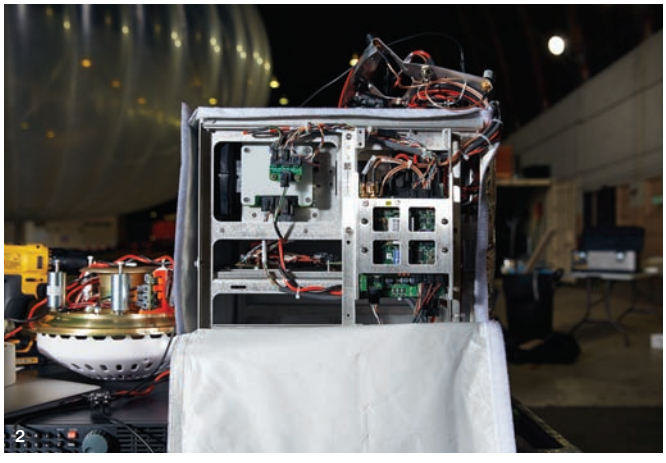

1 The helium balloons above are inflated to the size they reach in the stratosphere. The "ballonets" inside are filled with air or emptied to make the balloon fall or rise.

2 The 15-kilogram box carried by a Loon balloon has computers that act on commands from flight engineers, as well as equipment to transmit Internet connectivity to the ground below.

3 A balloon that was intentionally burst in a test is checked for flaws.

4 Loon balloons can be landed by carefully releasing helium, but they pack a parachute for emergencies.

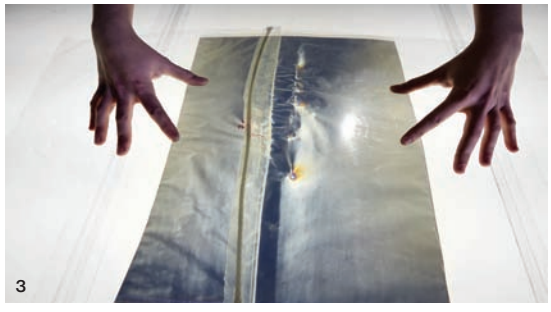

3

4

in Internet markets such as the United States. Getting billions more people online would provide a valuable new supply of eyeballs and personal data for ad targeting. That's one reason Project Loon will have competition: in 2014 Facebook bought a company that makes solar-powered drones so it can start its own airborne Internet project.

Google's planet-scale social-engineering project is much further along. In tests with major cellular carriers, the balloons have provided high-speed connections to people in isolated parts of Brazil, Australia, and New Zealand. Mike Cassidy, Project Loon's leader, says the technology is now sufficiently cheap and reliable for Google to start planning how to roll it out. By the end of 2015, he wants to have enough balloons in the air to test nearly continuous service in several parts of the Southern Hemisphere. Commercial deployment would follow: Google expects cellular providers to rent access to the balloons to expand their networks. Then the number of people in the world who still lack Internet access should start to shrink, fast.

Balloon revolution

"HARMLESS SCIENCE EXPERIMENT." That's what was written on the boxes carried by the balloons that the secretive Google X lab began to launch over California's Central Valley in 2012, along with a phone number and the promise of a reward for safe return. Inside the boxes was a modified office Wi-Fi router. The balloons were made by two seamsters hired from the fashion industry, from supplies bought at hardware stores.

Project Loon is now much less like a science project. In 2013, Google began working with a balloon manufacturer, Raven Aerostar, which expanded a factory and opened another to make the inflatable "envelope" for the balloons. That June, Google revealed the existence of the project and described its first small-scale field trials, in which Loon balloons provided Internet service to people in a rural area of New Zealand. In 2014, Project Loon focused on turning a functional but unwieldy prototype into technology that's ready to expand the world's communication networks.

Loon's leaders planned to buy their own space on the radio spectrum so their balloons could operate independently of existing wireless networks. But Google CEO Larry Page nixed that idea and said the balloons should instead be leased to wireless carriers, who could use the chunks of the airwaves they already own and put up ground antennas to link the balloons into their networks. That saved Google from spending billions on spectrum licenses and turned potential competitors into allies. "Nearly every telco we talk to wants to do it," says Cassidy.

Google has also made major improvements to its stratospheric craft. One of the most significant was developing a

way to accurately pilot balloons across thousands of miles without any form of propulsion. The stratosphere, which typically is used only by weather balloons and spy planes, is safely above clouds, storms, and commercial flights. But it has strong winds, sometimes exceeding 300 kilometers per hour. Providing reliable wireless service means being able to guarantee that there will always be a balloon within 40 kilometers.

Google solved that aviation problem by turning it into a computer problem. Winds blow in different directions and at different speeds in different layers of the stratosphere. Loon balloons exploit that by changing altitude. As a smaller balloon inside the main one inflates or deflates, they can rise or fall to seek out the winds that will send them where Google wants them to go. It's all directed by software in a Google data center that incorporates wind forecasts from the U.S. National Oceanic and Atmospheric Administration into a simulation of stratospheric airflow. "The idea is to find a way through the maze of the winds," says Johan Mathe, a software engineer working on Loon's navigation system. A fleet of balloons can be coördinated that way to ensure there is always one over any particular area.

The first version of this system sent new commands to Loon balloons once a day. It could find a way for a balloon launched over New Zealand, for example, to dawdle over land until prevailing winds pushed it east and over the Pacific Ocean. Then it would have the balloon ride the fastest winds possible for the 9,000-kilometer trip east to Chile. But that system could only get balloons within hundreds of kilometers of their intended target. For tests of Internet service in New Zealand and elsewhere, the company had to cheat, launching

Project Loon aims to change the economics of Internet access.

Loon balloons nearby to make sure they would be overhead. In late 2014, Google upgraded its balloon navigation system to give balloons fresh orders as frequently as every 15 minutes. They can now be steered with impressive accuracy over intercontinental distances. In early 2015, a balloon traveled 10,000 kilometers and got within 500 meters of its desired cell tower.

Google has also had to figure out how to make the balloons sturdier, so they can spend more time in the stratosphere. The longer they stay up, the lower the cost of operating the network. However, weight considerations mean a balloon's envelope must be delicate. Made from polyethylene plastic with the feel of a heavy-weight trash bag, the material is easily pierced with a fingertip, and a stray grain of grit in the factory can make a pinprick-size hole that will bring a balloon back to earth after less than two weeks.

Preventing those leaks is the work of a squad inside Project Loon that has doggedly chased down every possible cause and come up with preventive measures. These researchers have studied balloons retrieved from the stratosphere, pored over video footage of others inflated to bursting on the ground, and developed a "leak sniffer" to find tiny holes by detecting helium. The leak squad's findings have led to changes in the design of the balloon envelope, fluffier socks for factory workers who must step on the envelopes during production, and new machines to automate some manufacturing steps. Altogether, Google has introduced the first major changes the balloon industry has seen in decades, says Mahesh Krishnaswamy, who oversees manufacturing for Project Loon and previously worked on Apple's manufacturing operations. Those changes have paid off. In the summer of 2013, Loon balloons lasted only eight days before having to be brought down, says Krishnaswamy. Today balloons last on average over 100 days, with most exceeding that time in flight; a handful last as long as 130 days.

Google has also made many improvements to the design of the Loon balloons' payloads and electronics. But it still has problems left to solve. For example, Google needs to perfect a way of making radio or laser connections between balloons, so that they can pass data along in an aerial chain to connect areas far from any ground station.

But Cassidy says Project Loon's technology is already at a point where stratospheric Internet service can be tested at a global scale. In 2015 he aims to evaluate "quasi-continuous" service along a thin ribbon around the Southern Hemisphere. That ribbon is mostly ocean, but it will require a fleet of more than 100 Loon balloons circling the globe, says Cassidy. "Maybe 90 percent of the time," he says, "people in that ring will have at least one balloon overhead and be able to use it."

Good signals

"It was just for some minutes, but it was wonderful," says Silvana Pereira, a school principal in a rural area of northeastern Brazil. She's thinking back to an unusual geography class last summer in which pupils at Linoca Gayoso Castelo Branco School could use the Internet thanks to a Loon balloon drifting, invisibly, high overhead. Internet service is nonexistent in the area, but that day's lesson on Portugal was enhanced by

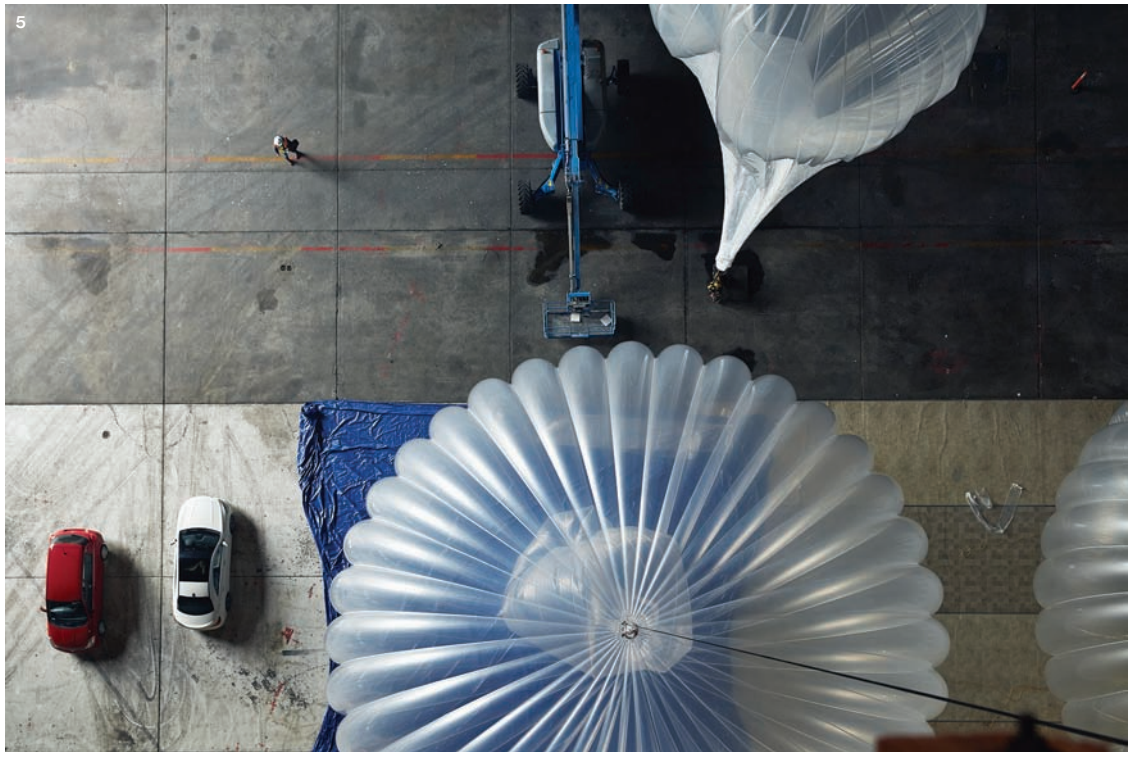

5 Among the upgrades Google is testing for its balloons (seen here from the rafters): using hydrogen, which is cheaper than helium, and having a motor move their solar panels to track the sun.

Wikipedia and online maps. "They were so involved that the 45 minutes of a regular class wouldn't be enough to satisfy their demand for knowledge," says Pereira.

Her school is only around 100 kilometers from a metro area of more than one million people, but its location is too poor and sparsely populated for Brazil's wireless carriers to invest in Internet infrastructure. Google's goal is for Project Loon to change those economics. It should be possible to operate one Loon balloon for just hundreds of dollars per day, Cassidy says, and each one should be able to serve a few thousand connections at any time. The company won't reveal how much it is spending to set all this up, or even how many people work on the project.

Cassidy is also confident that his balloons will be able to hold their own against Internet delivered by drones (both Google and Facebook are working on that) or satellites (an idea being pursued by SpaceX CEO Elon Musk). Those projects are less far along than Loon, and it's expensive to build and power drones or launch satellites. "For quite some time, balloons will have a big cost advantage," Cassidy says. (Google might be hedging its bets with more than just drones: it recently joined with Fidelity Investments to make a $1 billion investment in SpaceX.)

Technology is not the only thing keeping 4.3 billion people offline, though. For example, policies in India mandate that telecom companies provide coverage to poor as well as rich areas, but the government hasn't enforced the rules, says Sunil Abraham, executive director of the Centre for Internet and Society, a think tank in Bangalore. He is also wary of Project Loon because of the way Google and other Western Internet companies have operated in developing countries in recent years. They have cut deals with telecoms in India and other countries to make it free to access their websites, disadvantaging local competitors. "Anyone coming with deep pockets and new technology I would welcome," he says, but he adds that governments should fix up their patchy regulatory regimes first to ensure that everyone—not just Google and its partners—really does benefit.

Those working on Project Loon are confident the public good will be served. They seem as motivated by a desire to make people's lives better as by Loon's outlandish technology. Cassidy's voice wavers with emotion when he thinks back to seeing the delight of Pereira's pupils during their Internet-enabled geography lesson. "This is a way of changing the world," he says.

The Liquid Biopsy

Fast DNA-sequencing machines are leading to simple blood tests for cancer.

By Michael Standaert

Everything about China is big, including its cancer problem. In some wealthier cities, like Beijing, cancer is now believed to be the most frequent killer. Air pollution, high rates of smoking, and notorious "cancer villages" scarred by industrial pollution are increasing death rates around the country. Liver cancer in particular is four times as prevalent as it is in the West, in part because one in 14 people in China carry hepatitis B, which puts them at risk. Of all the people worldwide who die of cancer each year, some 27 percent are Chinese.

In December, I traveled by metro from Shenzhen to Hong Kong. There I had arranged to meet Dennis Lo, a doctor who has worked for nearly 20 years on a technique called the "liquid biopsy," which is meant to detect liver and other cancers very early—even before symptoms arise—by sequencing the DNA in a few drops of a person's blood.

Lo appeared fastidiously dressed as usual in a sharp blazer, a habit that called to mind formal dinners at the University of Oxford, where he studied in the 1980s. He is well known for having been the first to show that a fetus sheds bits of its DNA into the bloodstream of its mother. That finding, first made in 1997, has in recent years led to a much safer, simpler screening test for Down syndrome. By now more than one million pregnant women have been tested.

Today Lo is competing with labs around the world to repeat that scientific and commercial success by developing cancer screening tests based on a simple blood draw. That's possible because dying cancer cells also shed DNA into a person's blood. Early on, the amount is vanishingly small—and obscured by the healthy DNA that also

JAVIER JAEN

circulates. That makes it difficult to measure. But Lo says the objective is simple: an annual blood test that finds cancer while it's curable.

Cancers detected at an advanced stage, when they are spreading, remain largely untreatable. In the United States, early detection is behind medicine's most notable successes in applying technology to cut deaths from common cancers. Half of the steep decline in deaths from colorectal cancer is due to screening exams like colonoscopies.

Lo's hospital is involved in two of the largest studies anywhere to prove that DNA analysis can also act as a screening test. The researchers are following a thousand people with hepatitis B to see if the DNA test can spot liver tumors before an ultrasound can. An even larger study is on nasopharyngeal carcinoma, a cancer that starts in the upper part of the throat. It's rare elsewhere in the world, but in south China men have a one in 60 chance of contracting it in their lifetimes.

This cancer appears to be linked to eating salted fish, as well as to a genetic susceptibility among Chinese and to infection by the Epstein-Barr virus, the germ that causes mononucleosis. The role of the virus, says Lo, creates a special situation. The test he developed searches for easy-to-spot viral DNA that dying cancer cells release into a person's plasma.

The study involves 20,000 healthy middle-aged men recruited in Hong Kong, and it's halfway done. Among the first 10,000 men screened, the researchers picked up 17 cases of cancer—13 of those at stage I, the earliest kind. Nearly all these men have now beaten the cancer with radiation treatment. The typical survival rate is less than 70 percent if patients seek out a doctor only when they have the most advanced symptoms, like a mass in the neck. "They would normally be just walking on the street not knowing that there was a time bomb waiting to go off, and now we have alarmed them," says Lo. As he sees it, every man in south China could be screened. One private hospital in Hong Kong has started offering the test already. "We believe it will save lives," he says.

Lo's lab is now locked in a technology race with scientists at other institutions, including Johns Hopkins University, to see if these ideas can turn into a general-purpose test for nearly any cancer, not only those involving a virus. The approach relies on gene-sequencing machines, which rapidly decode millions of short fragments of DNA that are loose in the bloodstream. The results are compared with the reference map of the human genome. Researchers can then spot the specific patterns of rearranged DNA that are telltale signs of a tumor.

Lo showed me several older sequencing machines during a tour of his laboratory, located at the Chinese University of Hong Kong. He says that the next generation of DNA sequencers, some no larger than a cell phone, could allow routine screening for cancer to become less expensive and far more widely used. For the time being, the cost of the DNA test being tried out on people at risk for liver cancer is still too high for routine use. Lo notes that the fetal tests were similarly expensive at first but that prices have since declined to as little as $800. That's led to much wider use. "The same thing should happen [with cancer]," he says.

Building on the foundations put in place by doctors like Lo, commercial interest in liquid biopsies has recently started to explode. Eric Topol, a professor of genomics at the Scripps Research Institute, predicted this January that the technology, applied to cancer and other diseases, will become the "stethoscope for the next 200 years." Jay Flatley, CEO of Illumina, the San Diego company that builds fast gene-sequencing machines, told investors this year that the market for such tests could be worth at least $40 billion. Calling the technology "perhaps the most exciting breakthrough" in cancer diagnostics, he said his company would begin offering researchers a liquid-biopsy test kit to facilitate the search for signs of cancer.

In addition to screening for cancer, liquid biopsies could be a way to help people already fighting the disease. Doctors can pick a drug according to the specific DNA mutation driving a cancer forward. Tests to identify the mutation are sometimes done on tissue taken from a tumor, but a noninvasive blood test would be appropriate in more cases. Lo told me that 40 percent of Chinese lung cancer patients have a mutation in one gene, *EGFR*, that would make them eligible for new targeted drugs.

Cancer comes in many types, and Lo says that for each, researchers must methodically make their case that liquid biopsies can really save lives. He believes he's close with nasopharyngeal cancer. "If you can screen and prognosticate in very common cancer types, that is the time when it will go mainstream," he says.

Breakthrough
A blood test to catch cancer early.

Why It Matters
Cancer kills some eight million people a year around the world.

Key Players
- Dennis Lo, Chinese University of Hong Kong
- Illumina
- Bert Vogelstein, Johns Hopkins

High Tide for Seawater Desalination

The world is experiencing a huge wave of seawater desalination projects.

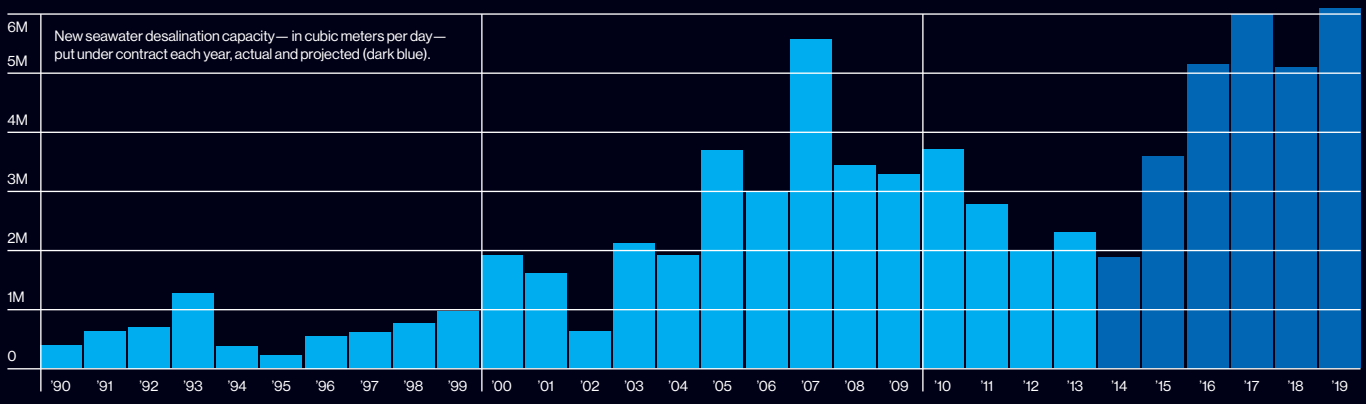

Fresh Water Is Drying Up

Colors map the ratio of how much fresh water people withdraw to how much is available.

Ratios of 0.4 or higher reflect high water stress imposed by the local population.

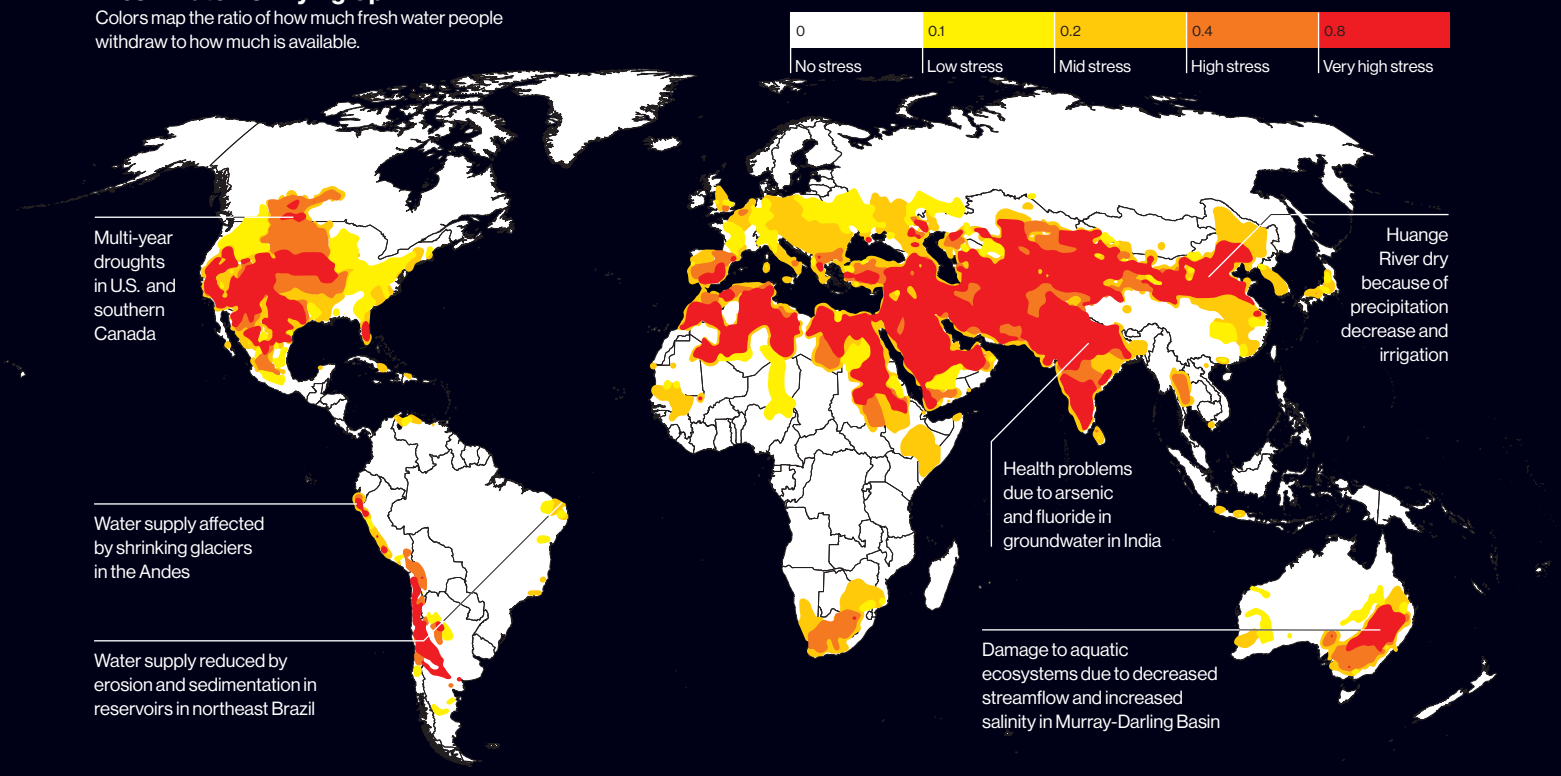

Israel's Turn to the Sea

Desalination will be providing 50 percent of drinkable water in Israel by 2016.

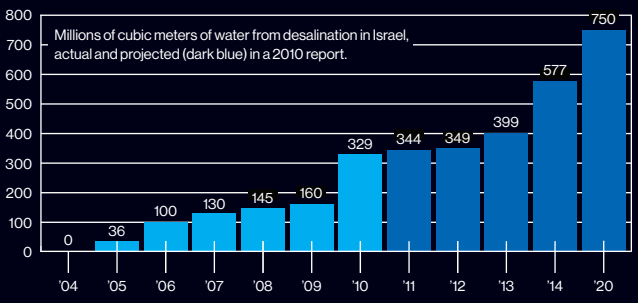

Water Supplies and Energy: No Free Drink

Seawater desalination still generally uses more energy than alternatives.

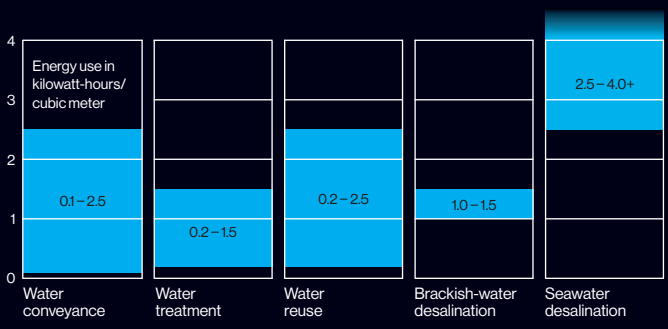

Megascale Desalination

The world's largest and cheapest reverse-osmosis desalination plant is up and running in Israel.

By David Talbot

Breakthrough
Demonstrating that seawater desalination can cost-effectively provide a substantial portion of a nation's water supply.

Why It Matters
The world's supplies of fresh water are inadequate to meet the needs of a growing population.

Key Players
- IDE Technologies
- Poseidon Water
- Desalitech
- Evoqua

On a Mediterranean beach 10 miles south of Tel Aviv, Israel, a vast new industrial facility hums around the clock. It is the world's largest modern seawater desalination plant, providing 20 percent of the water consumed by the country's households. Built for the Israeli government by Israel Desalination Enterprises, or IDE Technologies, at a cost of around $500 million, it uses a conventional desalination technology called reverse osmosis (RO). Thanks to a series of engineering and materials advances, however, it produces clean water from the sea cheaply and at a scale never before achieved.

Worldwide, some 700 million people don't have access to enough clean water. In 10 years the number is expected to explode to 1.8 billion. In many places, squeezing fresh water from the ocean might be the only viable way to increase the supply.

The new plant in Israel, called Sorek, was finished in late 2013 but is just now ramping up to its full capacity; it will produce 627,000 cubic meters of water daily, providing evidence that such large desalination facilities are practical. Indeed, desalinated seawater is now a mainstay of the Israeli water supply. Whereas in 2004 the country relied entirely on groundwater and rain, it now has four seawater desalination plants running; Sorek is the largest. Those plants account for 40 percent of Israel's water supply. By 2016, when additional plants will be running, some 50 percent of the country's water is expected to come from desalination.

The traditional criticism of reverse-osmosis technology is that it costs too much. The process uses a great deal of energy to force salt water against polymer membranes that have pores small enough to let fresh water through while holding salt ions back. However, Sorek will profitably sell water to the Israeli water authority for 58 U.S. cents per cubic meter (1,000 liters, or about what one person in Israel uses per week), which is a lower price than today's conventional desalination plants can manage. What's more, its energy consumption is among the lowest in the world for large-scale desalination plants.

The Sorek plant incorporates a number of engineering improvements that make it more efficient than previous RO facilities. It is the first large desalination plant to use pressure tubes that are 16 inches in diameter rather than eight inches. The payoff is that it needs only a fourth as much piping and other hardware, slashing costs. The plant also has highly efficient pumps and energy recovery devices. "This is indeed the cheapest water from seawater desalination produced in the world," says Raphael Semiat, a chemical engineer and desalination expert at the Israel Institute of Technology, or Technion, in Haifa. "We don't have to fight over water, like we did in the past."

Australia, Singapore, and several countries in the Persian Gulf are already heavy users of seawater desalination, and California is also starting to embrace the technology (see "Desalination Out of Desperation," January/February 2015). Smaller-scale RO technologies that are energy-efficient and relatively cheap could also be deployed widely in regions with particularly acute water problems—even far from the sea, where brackish underground water could be tapped.

Earlier in development are advanced membranes made of atom-thick sheets of carbon, which hold the promise of further cutting the energy needs of desalination plants.

ILLUSTRATION BY LUKE SHUMAN; DATA ON CAPACITY FROM GWI DESALDATA/IDA; MAP DATA FROM ALCAMO ET AL. (2003), IPCC; DATA ON ISRAEL FROM ISRAEL WATER AUTHORITY; DATA ON ENERGY USE FROM UNITED NATIONS

Apple Pay

A clever combination of technologies makes it faster and more secure to buy things with a wave of your phone.

By Robert D. Hof

SIMON LANDREIN

When Apple Pay was announced in September, Osama Bedier was unimpressed. A longtime PayPal executive who now runs a payment startup called Poynt, Bedier had spent more than two years leading Google's mobile wallet service, which lets people use their phones to pay for things at checkout counters. It used some of the same technologies as Apple Pay and failed to catch on widely. So despite Apple Pay's appealing promise—safe payment with just the press of a thumb on your iPhone—there was good reason to be skeptical of its chances, too.

Yet when Apple Pay launched just a few weeks later, Bedier was a convert. Poynt makes a new kind of payment terminal—one that retailers can use to accept Apple Pay—and the advent of the service helped send the company's orders soaring. "Now merchants have people walking in saying, 'Why can't I use Apple Pay?'" he says at Poynt's Palo Alto headquarters, whose lobby displays a 100-year-old National cash register, testament to the long history of payment technologies. Originally Bedier expected Poynt to sell 20,000 payment terminals in 2015, but after the launch of Apple Pay, he scrambled to find a new manufacturer in

Breakthrough
A service that makes it practical to use your smartphone as a wallet in everyday situations.

Why It Matters
Credit card fraud damages the economy by raising the costs of goods and services.

Key Players
- Apple
- Visa
- MasterCard
- Google


.10
00

Taiwan that could handle far greater demand. "Apple Pay will touch off a rush to mobile payment," he says.

Momentum for mobile payment technologies was building even before Apple Pay debuted last fall. Some 17 percent of all smartphone users reported making a point-of-sale payment with their phone in 2013, up from 6 percent in 2012, according to a U.S. Federal Reserve survey. In-person mobile payments in the United States more than doubled in 2014, to $3.7 billion, according to Forrester Research. Meanwhile, as services such as Uber and stores like Starbucks allow people to pay via mobile app, transactions that once brought out the wallet are disappearing into the phone, where they are faster and should be more secure. You can use your existing credit card accounts, but you never have to pull out the physical cards. "We know after people tap their phone to pay two or three times, they don't go back to their old behavior," says Ed McLaughlin, MasterCard's executive in charge of new payment technologies.

None of the individual technologies is novel, but Apple turned them into a service that is demonstrably easier than any other.

But even if Apple didn't invent mobile payments, it has significantly enhanced them. Just as Apple made it far easier to use a computer, listen to music, and communicate on the go, Apple Pay is all about doing the same for buying goods and services, online and off. Each financial innovation from the invention of money to the credit card reduced friction in commercial exchange and accelerated the pace of commerce. Apple Pay does that too: it marks the end of scrawling a signature, producing a driver's license, and other hassles that came with earlier forms of payment. It's also smoother than mobile services that came before it. Apple Pay works automatically when your phone is held up to the checkout terminal, with no need to open an app as you must to use Google Wallet or PayPal. Pressing your thumb to the phone eliminates the need to use a PIN, speeding the transaction. This is true no matter whether you're booking a room on Airbnb or buying sandwiches at Subway. It fuses the virtual and physical worlds of commerce in a way that no other payment system has done.

That doesn't mean most of us will be ditching our wallets and waving phones in every store in 2015—far from it. The $3.7 billion worth of mobile payments made in U.S. stores last year was just a drop in the $4 trillion bucket of consumer retail spending. Beyond that, an additional $12 trillion was spent on services. Apple Pay itself faces a raft of challenges, too, and not just from rival wallets offered by Google, PayPal, retailers, and wireless carriers. Currently only people with the new iPhone 6 can use Apple Pay in stores. It's officially available only in the United States for now, but 98 percent of U.S. stores lack the right checkout terminals to accept it. Finally, Apple Pay is far from replacing some of the things in a physical wallet—in particular, popular store rewards cards. Starbucks's app, which is a combination store locator, rewards card, and payment engine all in one, still accounts for the majority of all mobile payments in retail stores.

Still, Apple has done a lot of things right, suggesting that Apple Pay will turn out to be a milestone. Even if it is only a moderate success for Apple, it seems certain to be a driver of mobile payments in general. None of the individual technologies in it is novel, but the extent of Apple's control over both the software and the hardware in the iPhone—which exceeds what Google can do for Google Wallet even on Android phones—allowed it to combine those technologies into a service demonstrably easier to use than any other.

As a result, Apple is now cementing standards for the payment industry. Merchants have been debating whether bar codes or the radio technology near-field communication (NFC), for instance, should be the method that a phone uses to relay payment information when you wave it at a checkout terminal. Apple's choice to build NFC into iPhones means many stores will feel compelled to get terminals with NFC support if they want to maximize their appeal to millions of iPhone owners.

Likewise, Apple Pay is setting the pace in payment security, outdoing credit cards with multiple layers of protection (see "Tighter Security," next page). The phone doesn't store real card numbers, and even the merchant doesn't see them, let alone keep them in the databases that hackers routinely plunder. Each transaction generates a unique code that can be used only once. The capper: the payment is triggered with

Tighter Security

When you swipe a credit card at the checkout counter or buy something with your card online, you give the merchant your card number so the store can ask for approval from your card provider. The stores often keep those card numbers on their servers, where they repeatedly have been easy prey for criminals.

Apple Pay eliminates that exposure of your card number. When you sign up, you can use your phone's camera to take a picture of your card. Apple confirms the card with your bank, but then it deletes the photo, and the card number isn't stored on the phone or by Apple. Instead, Apple Pay creates an encrypted string of data called a device account number that stands in for your card. It gets stored on the phone in a special chip known as the Secure Element. The device account number can't be accessed by any applications on the phone other than Apple Pay. When it's time to buy something, the Secure Element coughs up the device account number and combines it with data about the transaction to create a unique code for that sale. A payment processor such as Visa or MasterCard is able to recognize the device account number and the unique code, and it uses them to approve or reject the transaction. The merchant never sees your actual card number.

Apple didn't invent this technology, and other payment services that use the wireless standard known as near-field communication also make use of secure elements. But Apple Pay goes a step further by combining these technologies with the iPhone's Touch ID fingerprint sensor, which is used to unlock the phone. That means you don't have to bother entering a PIN to confirm the transaction, but someone who steals your phone would be out of luck.

Touch ID, which responds only to the owner's fingerprint. This level of fraud protection is one reason banks representing 90 percent of U.S. consumer payments support Apple Pay, says Avin Arumugam, head of next-generation payment products at JP Morgan Chase.

Most of all, Apple's timing is impeccable. Card networks have set an October 2015 deadline for merchants to upgrade to terminals that can take credit cards with embedded chips for security—after that date, the merchants who don't upgrade will have to eat fraudulent charges. Most of those terminals they'll need to install will have NFC built in. Although that upgrade cycle will take years to reach most stores, Apple Pay could speed it up, says Keith Rabois, a former executive at PayPal and Square and an investor in several payment startups. "Apple Pay removes most of the barriers to adoption of mobile payment," he says.

Already, Apple Pay has taken off more quickly than Google Wallet or any other mobile payment system to date. "The time was ripe for Apple," says Jason Buechel, chief information officer at Whole Foods Market, where almost 2 percent of store sales were coming in through Apple Pay by mid-January. McDonald's said Apple Pay was accounting for half its mobile-phone transactions, and Walgreens's mobile payments doubled after Apple Pay debuted. Some 60 percent of customers used it on multiple days in November, using it three times as frequently as new PayPal customers used that system in the same time period, according to a study by the brokerage firm Investment Technology Group.

Apple stands to gain big if Apple Pay's momentum continues. Not from the 0.15 percent of each transaction that it charges card-issuing banks: those fees would bring in only $2.5 billion by 2017 even if the new system got an unexpectedly large 30 percent share of U.S. credit and debit card expenditures, according to one estimate by investor Carl Icahn. That's a tiny fraction of Apple's fiscal 2014 revenue of $183 billion. The bigger impact will be ensuring the iPhone's appeal. Once you're using Apple Pay every day, in addition to other Apple services like iCloud and iTunes, you may think thrice before switching to an Android.

For all the focus on Apple Pay in retail stores, its biggest opportunity in the next few years will probably be greasing payments for countless apps and services. When you take a ride with Uber, the payment happens almost invisibly, without friction. Rabois suggests that Apple Pay could bring that level of ease to thousands of on-demand services in transportation, food delivery, and more. Once people get used to making app payments with a touch, they'll start expecting to do the same everywhere else they can.

Brain Organoids

A new method for growing human brain cells could unlock the mysteries of dementia, mental illness, and other neurological disorders.

By Russ Juskalian

Photographs by Regina Huegli

Madeline Lancaster figured out a way to keep neurons growing in a dish until they develop characteristics of living human brains.

Breakthrough
Three-dimensional clusters of living neurons that can be grown in a lab from human stem cells.

Why It Matters
Researchers need new ways of understanding brain disorders and testing possible treatments.

Key Players
- Madeline Lancaster and Jürgen Knoblich, Institute of Molecular Biotechnology
- Rudolph Tanzi and Doo Yeon Kim, Massachusetts General Hospital

As Madeline Lancaster lifts a clear plastic dish into the light, roughly a dozen clumps of tissue the size of small baroque pearls bob in a peach-colored liquid. These are cerebral organoids, which possess certain features of a human brain in the first trimester of development—including lobes of cortex. The bundles of human tissue are not exactly "brains growing in a dish," as they're sometimes called. But they do open a new window into how neurons grow and function, and they could change our understanding of everything from basic brain activities to the causes of schizophrenia and autism.

Before it grows in one of Lancaster's dishes, a brain organoid begins as a single skin cell taken from an adult. With the right biochemical prodding, that cell can be turned into an induced pluripotent stem cell (the kind that can mature into one of several types of cells) and then into a neuron. This makes it possible to do things that were impossible before. Now scientists can directly see how networks of living human brain cells develop and function, and how they're affected by various drug compounds or genetic modifications. And because these mini-brains can be grown from a specific person's cells, organoids could serve as unprecedentedly accurate models for a wide range of diseases. What goes wrong, for example, in neurons derived directly from someone with Alzheimer's disease?

The prospect of finding answers to such questions is leading pharmaceutical companies and academic researchers to seek collaborations with Lancaster and Jürgen Knoblich, whose lab at the Institute of Molecular Biotechnology (IMBA) in Vienna, Austria, is where Lancaster developed the organoids as a postdoc. The first of these collaborations was an investigation of microcephaly, a disorder characterized by small brain size, with Andrew Jackson of the University of

1 Lancaster holds up organoids in a dish.

Organoids could be far more useful than animals in many experiments.

Edinburgh. Using cells derived from a patient with microcephaly, the team cultured organoids that shared characteristics with the patient's brain. Then the researchers replaced a defective protein associated with the disorder and were able to culture organoids that appeared partially cured.

This is just the beginning, says Lancaster. Researchers such as Rudolph Jaenisch at MIT and Guo-li Ming at Johns Hopkins are beginning to use brain organoids to investigate autism, schizophrenia, and epilepsy. What makes cerebral organoids particularly useful is that their growth mirrors aspects of human brain development. The cells divide, take on the characteristics of, say, the cerebellum, cluster together in layers, and start to look like the discrete three-dimensional structures of a brain. If something goes wrong along the way—which is observable as the organoids grow—scientists can look for potential causes, mechanisms, and even drug treatments.

The breakthrough in creating these organoids happened as part of a side project. Other researchers had grown neurons in a dish before, and like them, Lancaster started by using a flat plate to "play" with neural stem cells—the kind that form into neurons and other cells in the nervous system. Sometimes, she says, "I'd get neural stem cells that wouldn't really stay in 2-D, and they would kind of fall off the plate and they'd make 3-D clumps—and rather than ignoring them or throwing them away, I thought, 'Those are cool—let's see what happens if I let them keep growing.'" But there was a major challenge: how to keep the tissue at the center of the organoids fed without the benefit of veins. Lancaster's solution was to encapsulate each organoid in a matrix known to nurture cells, put a dozen of these blobs in a nutritious bath, and shake or spin it all to keep the organoids awash in cellular food.

Since publishing her method, Lancaster has pushed the brain tissue to further levels of complexity with neurons at later stages of development. The number of possible applications grows with each advance. Most tantalizing to Lancaster herself is the prospect that cerebral organoids might solve the deepest of mysteries: what happens in our brains to set us apart from other animals? "I'm mainly interested," she says, "in figuring out what it is that makes us human."

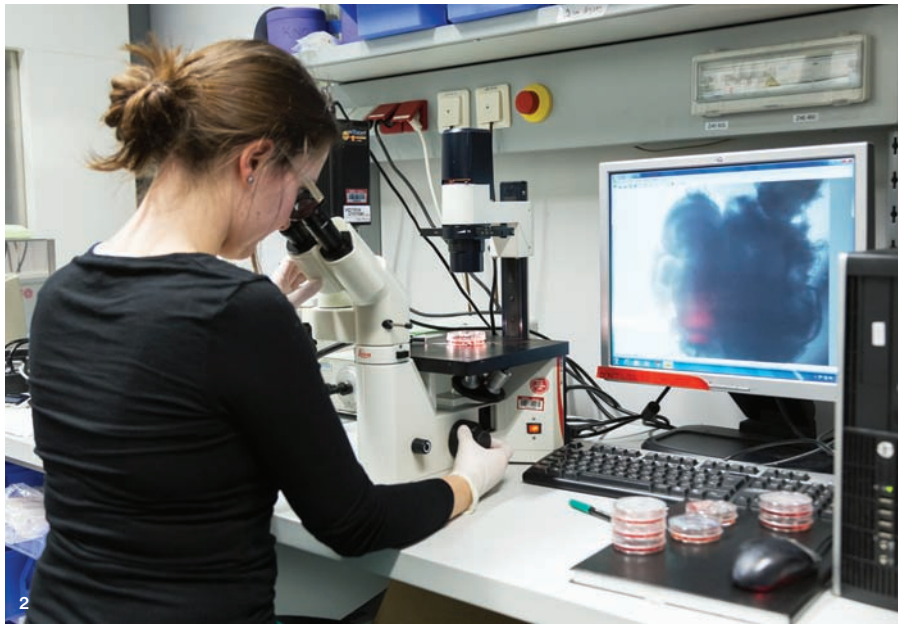

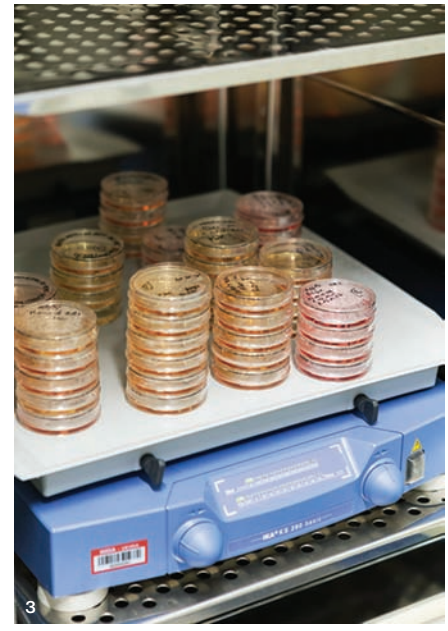

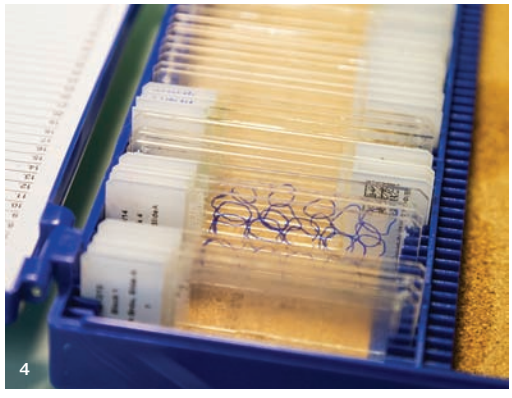

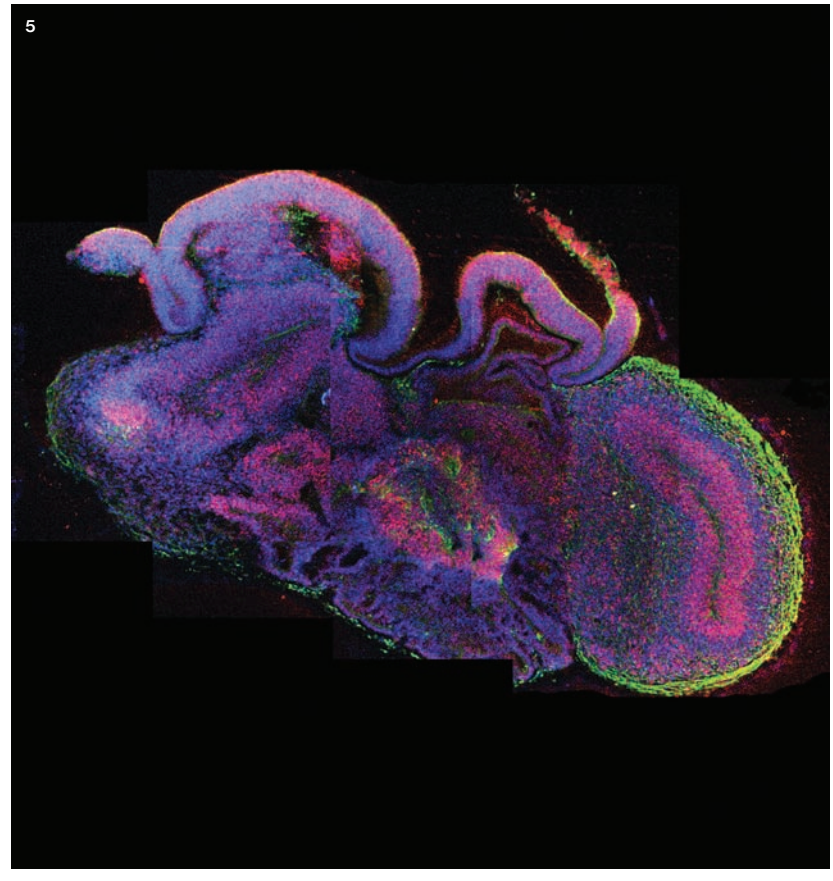

2 Magdalena Renner, a graduate student in the lab, examines organoids under a microscope.

3 A variety of organoids are kept alive on a shaker plate in an incubator.

4 Organoids cut into very thin sections have been put on slides for examination.

5 A stained section of an organoid is seen in close-up.

Photosynthesis Boost

The world's highest-production crops use a super-efficient form of photosynthesis. It's known as C4 photosynthesis because the first step is the formation of a four-carbon molecule. C3 photosynthesis, found in most plant species, starts with a three-carbon molecule.

Major C4 and C3 crops (annual production in metric tons)

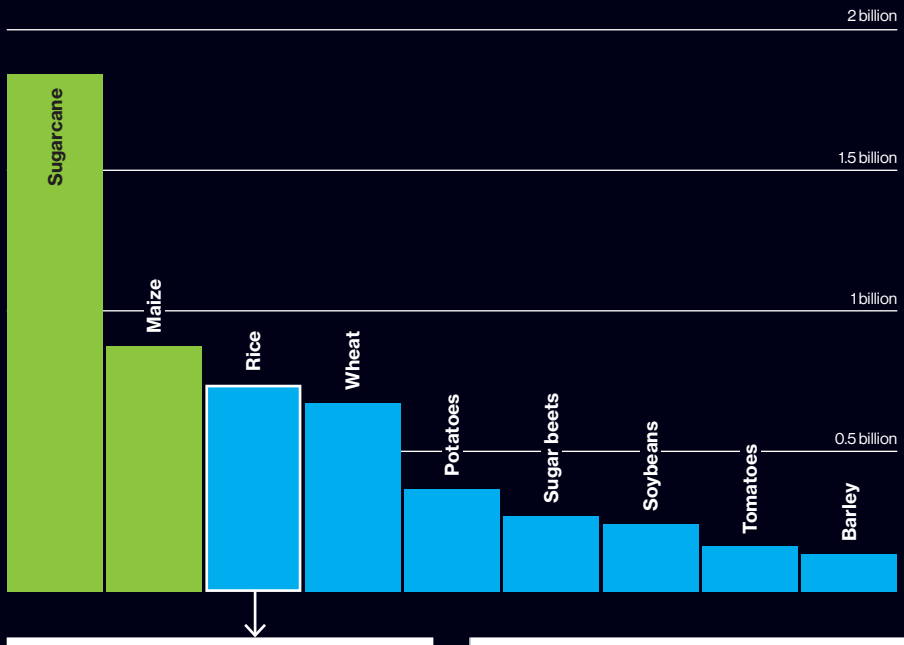

Carbon Concentrators

In C4 plants, a wreathlike arrangement of cells (lower image) helps concentrate carbon dioxide. A ring of mesophyll cells (green) captures the carbon dioxide, which is conveyed to an inner ring of bundle-sheath cells (orange). The arrangement is known as the Kranz anatomy, after the German word for wreath.

C3

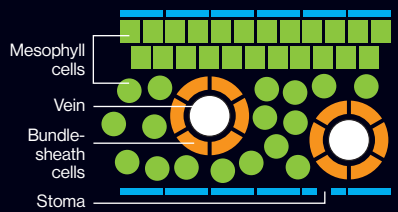

C4

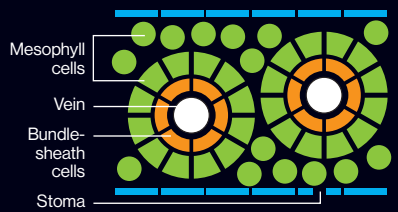

Rice Matters

Farmers are struggling to meet growing demand for rice, the staple for half of the world's population.

Rice provides 19% of global dietary energy

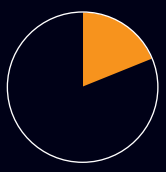

Plateauing yields

1990

Last year that average rice yields increased in California

33%

Percentage of rice-producing regions where yields have plateaued

Projected shortfall in rice production (in millions of tons)

2050 expected demand: 1,309

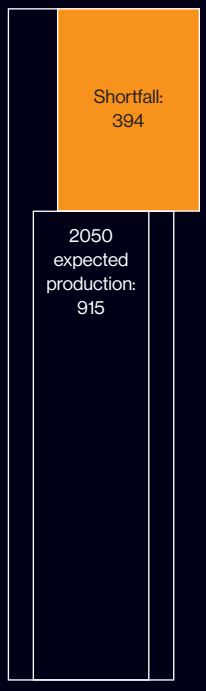

Efficient Farming

A unit of water goes further with C4 crops, producing far more food. In China, planting C4 rice could feed 50 percent more people per hectare.

Rice and corn grown with a given amount of water (the unit is a hectare covered to a depth of one millimeter with water)

People fed yearly in China by one harvest from one hectare of C3 vs. C4 rice

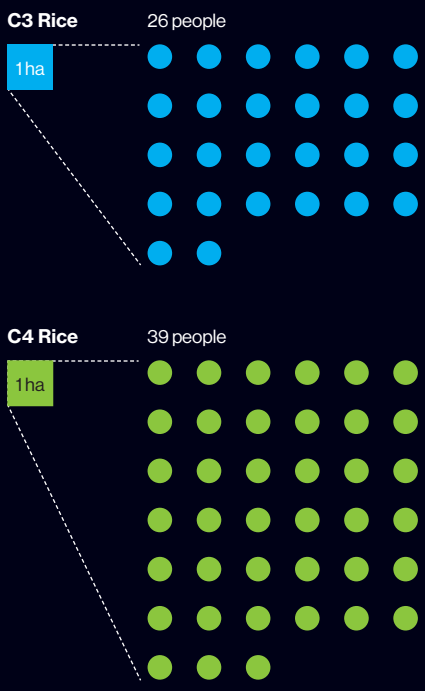

Supercharged Photosynthesis

Advanced genetic tools could help boost crop yields and feed billions more people.

By Kevin Bullis

Breakthrough
Engineering rice plants to extract energy from sunlight far more efficiently than they do now.

Why It Matters
Crop yields aren't increasing fast enough to keep up with demand from a growing population.

Key Players
- Paul Quick, International Rice Research Institute
- Daniel Voytas, University of Minnesota
- Julian Hibberd, University of Cambridge
- Susanne von Caemmerer, Australian National University

In December, geneticists announced that they'd engineered rice plants to carry out photosynthesis in a more efficient way—much as corn and many fast-growing weeds do. The advance, by a team of researchers in the Philippines and the United Kingdom, removes a big obstacle from scientists' efforts to dramatically increase the production of rice and, potentially, wheat. It comes at a time when yields of those two crops, which together feed nearly 40 percent of the world, are dangerously leveling off, making it increasingly difficult to meet rapidly growing food demand.

The supercharged process, called C4 photosynthesis, boosts plants' growth by capturing carbon dioxide and concentrating it in specialized cells in the leaves. That allows the photosynthetic process to operate much more efficiently. It's the reason corn and sugarcane grow so productively; if C4 rice ever comes about, it will tower over conventional rice within a few weeks of planting. Researchers calculate that engineering C4 photosynthesis into rice and wheat could increase yields per hectare by roughly 50 percent; alternatively, it would be possible to use far less water and fertilizer to produce the same amount of food.

The December results, achieved by Paul Quick at the International Rice Research Institute (IRRI) in the Philippines and Julian Hibberd, a professor at the University of Cambridge in the U.K., introduced five key C4 photosynthesis genes into a rice plant and showed that the plant can capture carbon dioxide via the same mechanism seen in plants with the supercharged form of photosynthesis. "It's the first time we've seen evidence of the C4 cycle in rice, so it's very exciting," says Thomas Brutnell, a researcher at the Danforth Plant Science Center in St. Louis. Brutnell is part of the C4 Rice Consortium headed by IRRI, which has funding from the Bill & Melinda Gates Foundation, but was not directly involved in the most recent breakthrough.

Despite the genetic changes, the altered rice plants still rely primarily on their usual form of photosynthesis. To get them to switch over completely, researchers need to engineer the plants to produce specialized cells in a precise arrangement: one set of cells to capture the carbon dioxide, surrounding another set of cells that concentrate it. That's the distinctive wreath anatomy found in the leaves of C4 plants. However, scientists still don't know all the genes involved in producing these cells and suspect that they could number in the dozens.

New genome editing methods that allow scientists to precisely modify parts of plant genomes could help solve the problem. Using conventional breeding to manipulate more than one or two genes is a "nightmare," Brutnell says, let alone trying to engineer a plant with dozens of gene changes. Genome editing could make it possible to change a large number of genes easily. Says Brutnell: "Now we have the toolbox to go after this."

It can be a decade or more before even simple crop modifications reach farmers, let alone changes as complex as reëngineering how plants carry out photosynthesis. But once scientists solve the C4 puzzle in a plant such as rice, they hope, the method can be extended to dramatically increase production of many other crops, including wheat, potatoes, tomatoes, apples, and soybeans.

ILLUSTRATION BY LUKE SHUMAN / SOURCES: FOOD AND AGRICULTURE ORGANIZATION OF THE UNITED NATIONS; CORNELL UNIVERSITY; INTERNATIONAL RICE RESEARCH INSTITUTE; PATRICIO GRASSINI, ET AL., NATURE COMMUNICATIONS; DEEPAK RAY ET AL., PLOSONE

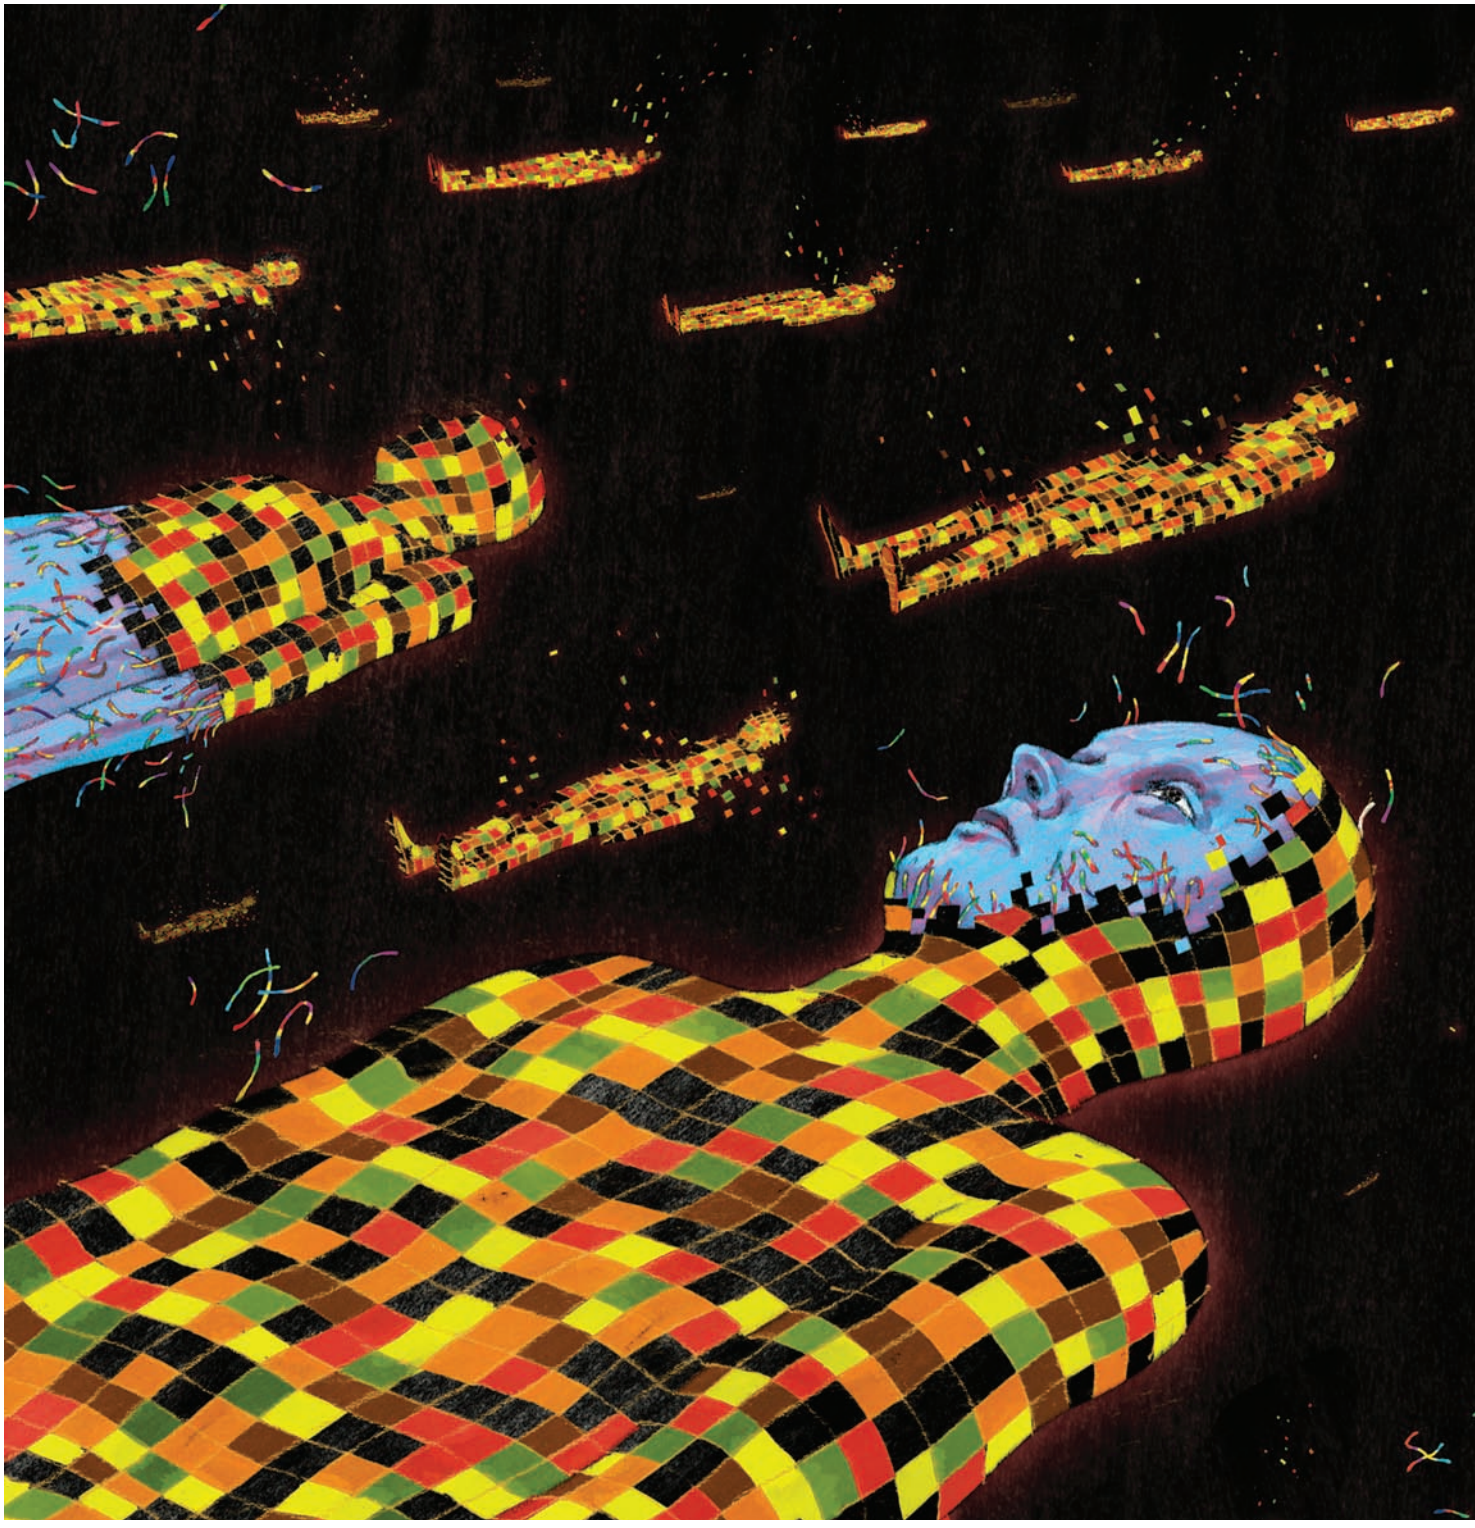

DADU SHIN

Internet of DNA

A global network of millions of genomes could be medicine's next great advance.

By Antonio Regalado

Noah is a six-year-old suffering from a disorder without a name. This year, his physicians will begin sending his genetic information across the Internet to see if there's anyone, anywhere, in the world like him.

A match could make a difference. Noah is developmentally delayed, uses a walker, speaks only a few words. And he's getting sicker. MRIs show that his cerebellum is shrinking. His DNA was analyzed by medical geneticists at the Children's Hospital of Eastern Ontario. Somewhere in the millions of As, Gs, Cs, and Ts is a misspelling, and maybe the clue to a treatment. But unless they find a second child with the same symptoms, and a similar DNA error, his doctors can't zero in on which mistake in Noah's genes is the crucial one.

In January, programmers in Toronto began testing a system for trading genetic information with other hospitals. These facilities, in locations including Miami, Baltimore, and Cambridge, U.K., also treat children with so-called Mendelian disorders, which are caused by a rare mutation in a single gene. The system, called MatchMaker Exchange, represents something new: a way to automate the comparison of DNA from sick people around the world.

One of the people behind this project is David Haussler, a bioinformatics expert based at the University of California, Santa Cruz. The problem Haussler is grappling with now is that genome sequencing is largely detached from our greatest tool for sharing information: the Internet. That's unfortunate because more than 200,000 people have already had their genomes sequenced, a number certain to rise into the millions in years ahead. The next era of medicine depends on

Breakthrough
Technical standards that let DNA databases communicate.

Why It Matters
Your medical treatment could benefit from the experiences of millions of others.

Key Players
- Global Alliance for Genomics and Health
- Google
- Personal Genome Project

large-scale comparisons of these genomes, a task for which he thinks scientists are poorly prepared. "I can use my credit card anywhere in the world, but biomedical data just isn't on the Internet," he says. "It's all incomplete and locked down." Genomes often get moved around in hard drives and delivered by FedEx trucks.

Haussler is a founder and one of the technical leaders of the Global Alliance for Genomics and Health, a nonprofit organization formed in 2013 that compares itself to the W3C, the standards organization devoted to making sure the Web functions correctly. Also known by its unwieldy acronym, GA4GH, it's gained a large membership, including major technology companies like Google. Its products so far include protocols, application programming interfaces (APIs), and improved file formats for moving DNA around the Web. But the real problems it is solving are mostly not technical. Instead, they are sociological: scientists are reluctant to share genetic data, and because of privacy rules, it's considered legally risky to put people's genomes on the Internet.

The unfolding calamity in genomics is that a great deal of life-saving information, though already collected, is inaccessible.

But pressure is building to use technology to study many, many genomes at once and begin to compare that genetic information with medical records. That is because scientists think they'll need to sort through a million genomes or more to solve cases—like Noah's—that could involve a single rogue DNA letter, or to make discoveries about the genetics of common diseases that involve a complex combination of genes. No single academic center currently has access to information that extensive, or the financial means to assemble it.

Haussler and others at the alliance are betting that part of the solution is a peer-to-peer computer network that can unite widely dispersed data. Their standards, for instance, would permit a researcher to send queries to other hospitals, which could choose what level of information they were willing to share and with whom. This control could ease privacy concerns. Adding a new level of complexity, the APIs could also call on databases to perform calculations—say, to reanalyze the genomes they store—and return answers.

The day I met Haussler, he was wearing a faded Hawaiian shirt and taking meetings on a plastic lawn chair by a hotel pool in San Diego. Both of us were there to attend one of the world's largest annual gatherings of geneticists. He told me he was worried that genomics was drifting away from the open approach that had made the genome project so powerful. If people's DNA data is made more widely accessible, Haussler hopes, medicine may benefit from the same kind of "network effect" that's propelled so many commercial aspects of the Web. The alternative is that this vital information will end up marooned in something like the disastrous hodgepodge of hospital record systems in the United States, few of which can share information.

One argument for quick action is that the amount of genome data is exploding. The largest labs can now sequence human genomes to a high polish at the pace of two per hour. (The first genome took about 13 years.) Back-of-the-envelope calculations suggest that fast machines for DNA sequencing will be capable of producing 85 petabytes of data this year worldwide, twice that much in 2019, and so on. For comparison, all the master copies of movies held by Netflix take up 2.6 petabytes of storage.

"This is a technical question," says Adam Berrey, CEO of Curoverse, a Boston startup that is using the alliance's standards in developing open-source software for hospitals. "You have what will be exabytes of data around the world that nobody wants to move. So how do you query it all together, at once? The answer is instead of moving the data around, you move the questions around. No industry does that. It's an insanely hard problem, but it has the potential to be transformative to human life."

Today scientists are broadly engaged in what is, in effect, a project to document every variation in every human gene and determine what the consequences of those differences are. Individual human beings differ at about three million DNA positions, or one in every 1,000 genetic letters. Most of these differences don't matter, but the rest explain many things that do: heartbreaking disorders like Noah's, for example, or a higher than average chance of developing glaucoma.

So imagine that in the near future, you had the bad luck to develop cancer. A doctor might order DNA tests on your tumor, knowing that every cancer is propelled by specific mutations. If it were feasible to look up the experience of

everyone else who shared your tumor's particular mutations, as well as what drugs those people took and how long they lived, that doctor might have a good idea of how to treat you. The unfolding calamity in genomics is that a great deal of this life-saving information, though already collected, is inaccessible. "The limiting factor is not the technology," says David Shaywitz, chief medical officer of DNAnexus, a bioinformatics company that hosts several large collections of gene data. "It's whether people are willing."

Last summer Haussler's alliance launched a basic search engine for DNA, which it calls Beacon. Currently, Beacon searches through about 20 databases of human genomes that were previously made public and have implemented the alliance's protocols. Beacon offers only yes-or-no answers to a single type of question. You can ask, for instance, "Do any of your genomes have a T at position 1,520,301 on chromosome 1?" "It's really just the most basic question there is: have you ever seen this variant?" says Haussler. "Because if you did see something new, you might want to know, is this the first patient in the world that has this?" Beacon is already able to access the DNA of thousands of people, including hundreds of genomes put online by Google.

One of the cofounders of the Global Alliance is David Altshuler, who is now head of science at Vertex Pharmaceuticals but until recently was deputy chief of the MIT-Harvard Broad Institute, one of the largest academic DNA-sequencing centers in the United States. The day I visited Altshuler in his Broad office, his whiteboard was covered with diagrams showing genetic inheritance in families, as well the word "Napster" written in large blue letters—a reference to the famously disruptive music-sharing service of the 1990s.

Altshuler has his own reasons for wanting to connect massive amounts of genetic data. As an academic researcher, he hunted for the genetic causes of common diseases like diabetes. That work was carried out by comparing the DNA of afflicted and unafflicted people, trying to spot the differences that come up most often. After burning through countless research grants this way, geneticists realized there would be no easy answers, no common "diabetes genes" or "depression genes." It turns out that common diseases aren't caused by single, smoking-gun defects. Instead, a person's risk, scientists have learned, is determined by a combination of hundreds, if not tens of thousands, of rare variations in the DNA code.

That's created a huge statistical headache. Last July, in a report listing 300 authors, Broad looked at the genes of 36,989 people with schizophrenia. Even though schizophrenia is highly heritable, the 108 gene regions identified by the scientists explained only a small percentage of a person's risk for the disease. Altshuler believes that big gene studies are still a good way to "crack" these illnesses, but he thinks it will probably take millions of genomes to do it.

The way the math works out, sharing data no longer looks optional, whether researchers are trying to unravel the causes of common diseases or ultra-rare ones. "There's going to be an enormous change in how science is done, and it's only because the signal-to-noise ratio necessitates it," says Arthur Toga, a researcher who leads a consortium studying the science of Alzheimer's at the University of Southern California. "You can't get your result with just 10,000 patients—you are going to need more. Scientists will share now because they have to."

Privacy, of course, is an obstacle to sharing. People's DNA data is protected because it can identify them, like a fingerprint—and their medical records are private too. Some countries don't permit personal information to be exported for research. But Haussler thinks a peer-to-peer network can sidestep some of these worries, since the data won't move and access to it can be gated. More than half of Europeans and Americans say they're comfortable with the idea of sharing their genomes, and some researchers believe patient consent forms should be dynamic, a bit like Facebook's privacy controls, letting individuals decide what they'll share and with whom—and then change their minds. "Our members want to be the ones to decide, but they aren't that worried about privacy. They're sick," says Sharon Terry, head of the Genetic Alliance, a large patient advocacy organization.

The risk of not getting data sharing right is that the genome revolution could sputter. Some researchers say they are seeing signs that it's happening already. Kym Boycott, head of the research team that sequenced Noah's genome, says that when the group adopted sequencing as a research tool in 2010, it met with immediate success. Over two years, between 2011 and 2013, a network of Canadian geneticists uncovered the precise molecular causes of 146 conditions, solving 55 percent of their undiagnosed cases.

But the success rate appears to be tailing off, says Boycott. Now it's the tougher cases like Noah's that are left, and they are getting solved only half as often as the others. "We don't have two patients with the same thing anymore. That's why we need the exchange," she says. "We need more patients and systematic sharing to get the [success rate] back up." In late January, when I asked if MatchMaker Exchange had yielded any matches yet, she demurred, saying that it could be a matter of weeks before the software was fully operational. As for Noah, she said, "We are still waiting to sort him out. It's important for this little guy."

Northern Ireland's Cybersecurity Cluster

Belfast's growing reputation as a technology hotspot has been further enhanced by the addition of three U.S. cybersecurity leaders joining an already strong cluster of indigenous and international companies.

Most recently setting up shop in the region is U.S. Web security specialist WhiteHat Security. The Santa Clara, California-based company will support its growth in Europe and beyond from its new Threat Research Centre (TRC) in Belfast, which will employ 50 security engineers.

"We considered a number of locations for our new TRC but determined that Belfast provided the right combination of talent and location that would help to serve our needs as well as the needs of our customers," said WhiteHat Security interim CEO and founder, Jeremiah Grossman.

Proofpoint, a Sunnyvale, California-based security-as-a-service company, will also expand its operation in Belfast, adding 90 new jobs. Proofpoint's Belfast operation will serve as the Research and Development and Operations Centre for Proofpoint Essentials, supporting the company's customer base in Europe, the Middle East, and Africa.

"We have been very impressed with the technology ecosystem in Northern Ireland, which shares many of the characteristics of Silicon Valley, including access to world-class cybersecurity research at Queen's University Belfast," said Proofpoint CEO Gary Steele.

Another recent arrival in Northern Ireland is Boston-based Rapid7. One of the United States' fastest-growing software-security companies, Rapid7 provides customers with insight into the security state of their IT, data assets, and users across virtual, mobile, private-cloud and public-cloud networks. Rapid7's new 75-employee development center in Belfast will help the company scale its operations and reach a wider customer base in the United States, Europe, and Asia.

"We were impressed with the excellent business environment, particularly the close working relationship between government, industry and academia. Of note was the work the universities are doing in IT security," said Rapid7 CEO Corey Thomas..

These international cybersecurity companies are finding that operating in Northern Ireland allows them to benefit from state-of-the-art research and development capability and to tap into an exceptional pool of talent.

The region is home to the Centre for Secure Information Technologies (CSIT) at Queen's University Belfast. CSIT has been recognized as the U.K.'s leading research center for cybersecurity by the U.K. Government Communications Headquarters.

The commercially focused center concentrates on secure ubiquitous networking, device authentication, security analytics, and autonomous sensor security.

Multinational companies such as McAfee, Q1 Labs, Tyco, Cisco, Chemring, IBM, Allstate, and BAE Systems have become members of CSIT, allowing them to influence intellectual property produced there.

In September 2014, CSIT and Queen's University Belfast launched a new master's degree program in cybersecurity with the goal of ensuring that Belfast continues to provide professionals with the skills needed to serve this fast-moving industry.

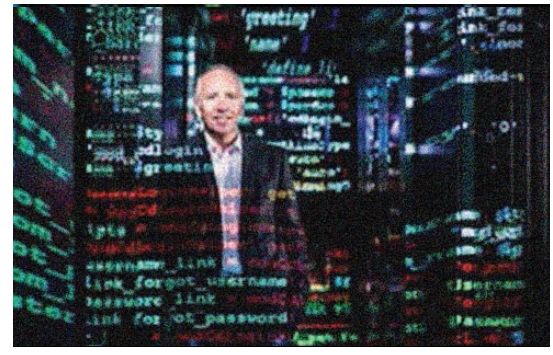

Above: Dr Godfrey Gaston, Director, CSIT, the UK's Innovation and Knowledge Centre (IKC) for secure information technologies

In addition, Northern Ireland also offers strong communication and travel advantages. The region benefits from a super-speed telecommunications link to North America. Three airports provide gateways to London and the world, and direct flights to New York are available.

You can meet representatives from Invest NI, Northern Ireland security companies, and CSIT at the RSA Conference in San Francisco in April 2015 and at InfoSecurity Europe in London in June 2015.

To find out more about what Northern Ireland has to offer your business, visit investni.com

MIT Technology Review
BUSINESS REPORT

CONTENTS

The Future of Money

From Apple Pay to Bitcoin, payment technology is changing. Some of the big winners in this "revolution" may be the old guard.

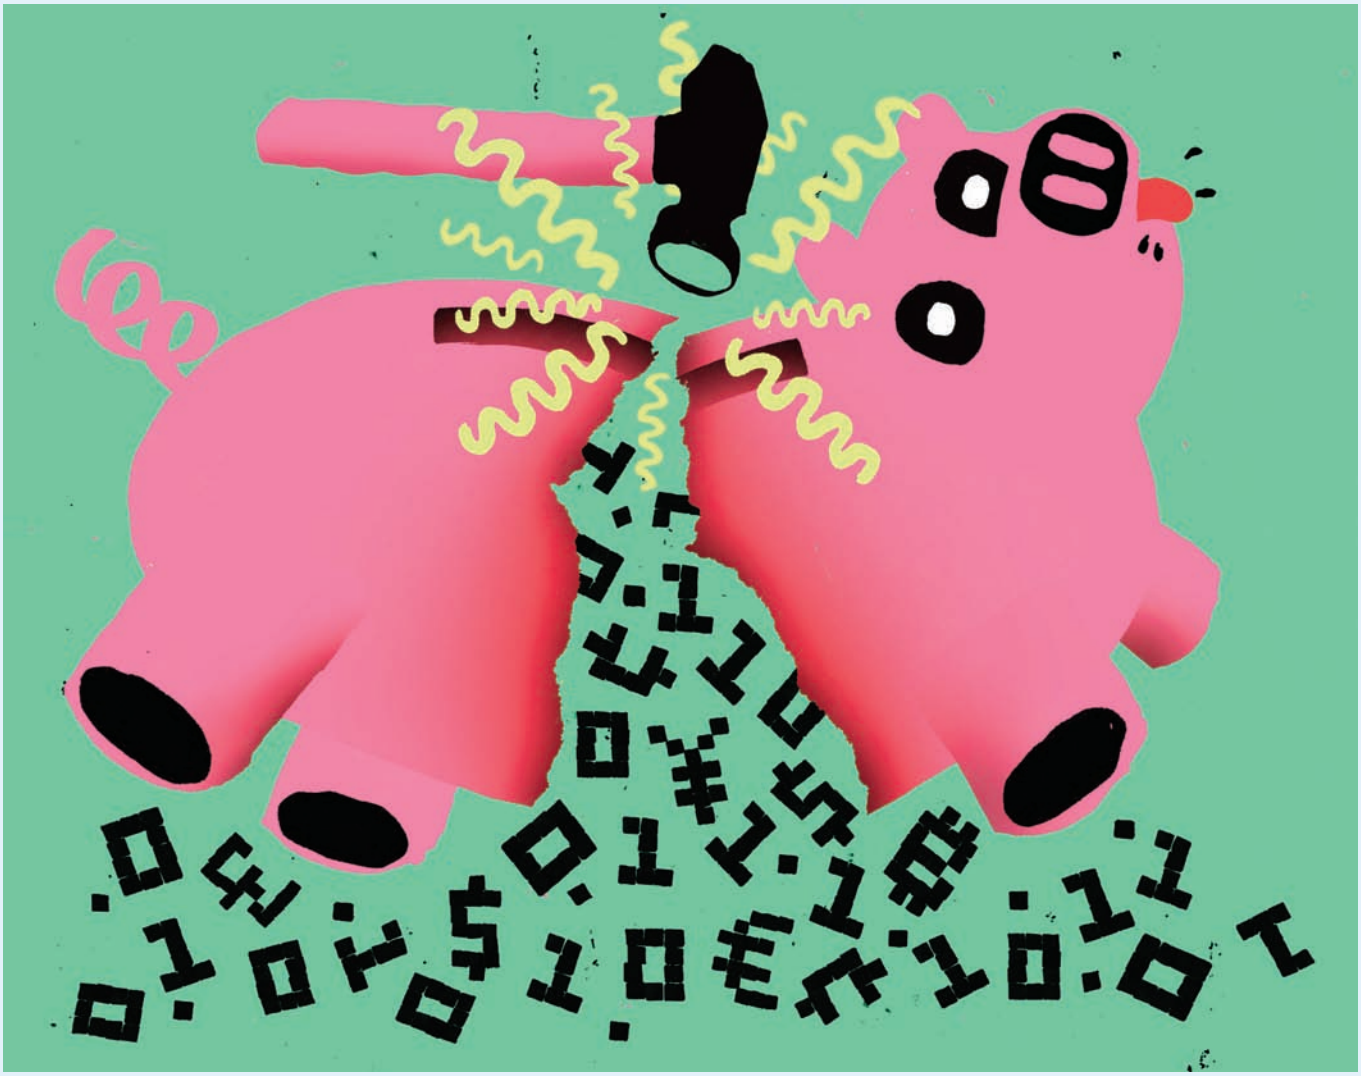

DANIEL ZENDER

The Big Question

Technology Repaints the Payment Landscape

As technology-driven payment ideas give cash a run for its money, the big winners could be established banks and credit card companies.

● In developed economies, money has been digitizing for decades. Few Westerners touch a paycheck anymore. Through direct deposit, digital money is transferred electronically from our employer to our bank account every pay period. A similar process moves contributions into our 401(k) accounts or zaps money over to pay the rent, the utility bill, a student loan, or any other expense.

Yet it remains a cash-based world, with 85 percent of consumer transactions worldwide done with bills and coins. While some countries, like Singapore and the Netherlands, now use cash in a minority of payments, consumers in such diverse economies as India, Mexico, Italy, and Taiwan still execute more than 90 percent of transactions with cash, according to research by MasterCard Advisors. Even in the United States, they find, cash accounts for 55 percent of payments. New technologies, including digital wallets, cryptocurrencies, and mobile peer-to-peer payments, aim to tip that balance. They're accelerating the move away from cash in countries where alternatives to banks and credit cards are well established, and they're doing the same in developing economies.

Which technologies and companies are likely to lead this transformation is the big question for this Business Report.

One way to look at these new technologies is through their relationship with the long-established payment services. Some technologies, including the mobile wallets Apple Pay and LoopPay, run on top of the existing payment networks owned and operated by banks and credit card companies. The new technologies are designed to make those established systems faster, more convenient, or more secure, and to convert transactions now being done in cash. A different group of technologies would replace the established systems with new ones, fundamentally challenging the vast industry that executes, guarantees, and tracks payments. Among them: Venmo, a person-to-person payment app and social network that processes $3 billion of payments a year, and Dwolla, an upstart in Iowa looking to cut the payment-processing revenue enjoyed by Visa and other networks.

As technology drives a shift in how we buy things, the revenue that the payments industry extracts could grow to more than $2 trillion a year by 2023, double the 2013 figure, the Boston Consulting Group predicts. Much of that increase will come from a reduction in cash payments in developing countries. But across the globe, BCG predicts a time of "disruption and opportunity" driven by digital technologies that will require the existing credit card system to prove that it's better than its new competition.

"The smartphone is the catalyst for a lot of change in this industry," says Dana Stalder, a venture capitalist with Matrix Partners and a former eBay and PayPal executive now on the board of Poynt, which recently introduced a smart credit card terminal. Venture capitalists invested over $2 billion in payment technology firms between January 2013 and June 2014, according to the data tracking firm CB Insights.

However, established players, especially the banks and credit card companies that handle most noncash payments today, have, if anything, seen their positions strengthened by recent developments. A good example is the high-profile launch of Apple Pay. Unlike earlier technologies like Google Wallet and PayPal, Apple Pay makes no attempt to supplant the Visas and Bank of Americas of the world. Look at your digital wallet in Apple Pay and you see a version of exactly the same card that's in the wallet in your pocket. The digital wallet LoopPay, which can be used in many more terminals than Apple Pay because it uses a simple, widely compatible copper loop technology to simulate the coding in your credit card's magnetic strip, similarly relies on the existing credit card system.

$2 billion
VC investment in payment tech

"Think about the infrastructure and how long it took to create that," says LoopPay CEO Will Wang Graylin. "It's very difficult to change merchant behavior."

Innovation in payments might be especially likely to take hold in the devel-

The Surprising Persistence of Cash
Estimated percentage of consumer payment transactions done with noncash methods, including credit cards, checks, and electronic payments.

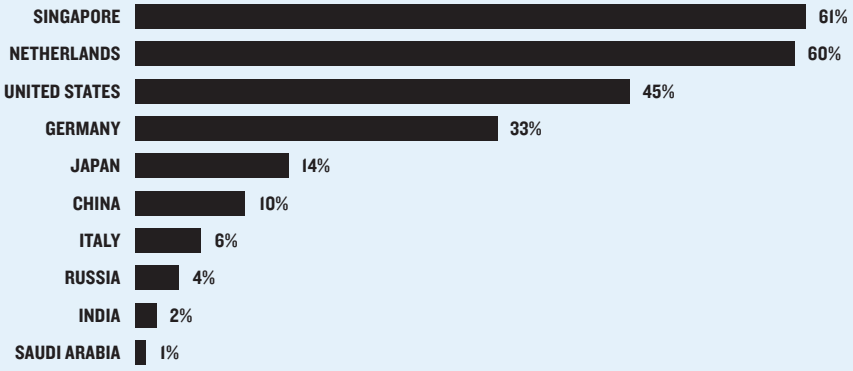

SOURCE: MASTERCARD ADVISORS

oping world, where cash is still king. Leapfrogging ATMs and checks the same way they have skipped over landlines and cable, whole chunks of population are moving straight from cash to mobile money. M-Pesa, which has become a force in Kenya and Tanzania, has turned money into a cellular currency that can be converted into airtime or used to pay for things. Today, some 60 percent of Kenyan adults have used a mobile phone to receive or send payments.

What could derail the boom in payment technologies?

Security concerns. The consulting firm Accenture recently surveyed 4,000 consumers in North America and found that while more people expect to use mobile payments, 57 percent of respondents were concerned about the security of such transactions. That's up from 45 percent two years ago.

New approaches could help. Apple Pay, Google Wallet, and others utilize a system that creates a one-time digital token for each transaction and sends that, rather than a customer's credit card information, through the system.

Innovations like this show that mobile payments—even if they don't lead to a radical shake-up—are improving a global payment ecosystem that's long been due for an upgrade. *—Nanette Byrnes*

Case Study

Alipay Leads a Digital Finance Revolution in China

What started as a service to help customers buy goods from Alibaba retailers has grown into a serious finance business all its own.

● Not many years ago, Jane Yang, a 26-year-old civil servant in Beijing, paid her landlord in three-month installments with a stack of 100-yuan notes. To pay her utilities—water, electricity, and home Internet—she went to three separate banks, where she handed cash to a teller. The process was "very time-consuming and irritating," she remembers. Even as skyscrapers and gleaming shopping malls cropped up around China's capital, most middle-class residents had never seen or used a simple checkbook.

Today she uses the Alipay app, China's most popular online payment service, on her smartphone to transfer money directly to her landlord's account. She pays for her utilities and her mobile-phone account through Alipay as well. Yang even keeps savings in her Alipay Yu'ebao money market account, where money accrues higher interest than it does in a traditional bank account. Yang hadn't set out to deliberately overhaul her financial habits, but new mobile technology, she says, "made so it easy."

As of October 2014, Alipay had more than 300 million registered users in China (and 17 million overseas), according to the most recent figures the company has made public. Many, like Yang, originally set up accounts in order to shop at parent company Alibaba's wildly popular retail sites Taobao.com and Tmall.com, where everything from designer clothes to pet food is for sale. The Alipay payment system works much like PayPal, except that funds are held in escrow and are released when the goods arrive in satisfactory fashion. In a society where consumers have learned to be wary of false advertising and fake products, Alipay's escrow system helped ease consumer fears—and gave Alibaba's retail sites a crucial early advantage over rival eBay.

Today Alibaba's sites sell $300 billion worth of goods annually, dwarfing sales on eBay and Amazon combined. The company, which unlike Amazon doesn't actually stock and sell the merchandise on its sites, held its initial public offering on the New York Stock Exchange in September, raising $25 billion, the largest debut ever.

"[Alibaba and Alipay] are integral to each other's success. But I wouldn't be surprised if, in the long term, Alipay turns out to be the more important business."
—Ben Cavender, China Market Research Group

Alibaba and Alipay, which has been incorporated as a separate company since 2011, helped drive the very rapid expansion of online sales in China—now the world's second-largest "e-tail" market. McKinsey Global Institute estimates that by 2020, Chinese e-tailing could generate as much as $650 billion in sales, and China's market "will equal that of the United States, Japan, the United Kingdom, Germany, and France combined today."

As much as Alibaba has driven China's booming e-commerce market, it's possible that Alipay will ultimately have the bigger impact on the Chinese economy. Alibaba and Alipay "are integral to each other's success," says Ben Cavender, principal at China Market Research Group in Shanghai. "But I wouldn't be surprised if, in the long term, Alipay turns out to be the more important business—it's so flexible and has so many potential uses."

Alipay debuted as a simple e-payment system, but it's now a destination app (and website) in its own right. In addition to easing consumers into online shopping, Alipay, with its huge built-in user base, has recently made a range of financial services available to people who previously lacked easy access to money market accounts, small business loans, and tools for making payments.

As the Chinese fast become accustomed to banking on their phones, Alipay faces new competition from alternatives like Tencent's Weixin Wallet function, which enables mobile payments.

"In China today, it's technology companies that are driving innovation in

mobile payments," says Zennon Kapron, a financial technology and digital currency expert in Shanghai. Traditional banks and the government are not the main actors.

Between July 2013 and June 30, 2014, Alipay handled $778 billion (4.8 trillion yuan) in transactions, according to the company. It is able to process more than 10 billion transactions per day. During the popular "Singles' Day" annual sale—which is like Black Friday in the U.S. but on overdrive—Alipay handled up to 2.85 million transactions per minute, and 54 percent of its transactions are made via mobile device.

Alipay's back-end technology is similar to that of PayPal, Kapron says, but on the front end the user's experience is quite different. PayPal is best known as a payment option, a screen you may reach at the end of a transaction on a retail website, but with Alipay, customers can go directly to its app and website to make payments, check their investments, or buy movie or plane tickets. The presentation, says Kapron, is suited "for people who have quickly gotten used to using their mobile phones for everything."

With these new mobile payment technologies, China has leapfrogged both checkbooks and desktop banking. Jane Yang, for example, went straight from paying rent in cash to paying via Alipay. According to PricewaterhouseCoopers, 79 percent of Chinese consumers surveyed said they were happy to receive coupons via their mobile devices, versus just 53 percent globally. And 55 percent of Chinese consumers said they expected their phone to be the main way they made purchases in the future, versus 29 percent globally.

This is a remarkable turnaround for a country that for years seemed to be stuck in a far earlier, low-tech era of consumer financial services. "The banks did nothing to make customer service easy," says Cavender, who notes that for many years paying a credit card bill required standing in line at a bank. It could not be done through the mail or online. Only those who had significant funds to invest and lived near large bank branches had easy access to wealth management options.

These changes coincide with rising overall incomes in China, and with the government's desire to build a more consumer-based society, observes Tjun Tang, senior partner and managing director at Boston Consulting Group's Hong Kong office. "In uptake of digital finance, China is probably leading the world right now," he says. *—Christina Larson*

Hacking Payments

Attacks that come in many forms are yielding more records and doing real damage.

Incidents

Some of the most notable large-scale breaches of card information, by day of announcement.

Impact

Records, in millions, exposed in incidents involving stolen payment card data in retail and finance. Includes breaches publicly disclosed in English-language media reports, government sources, and online forums.

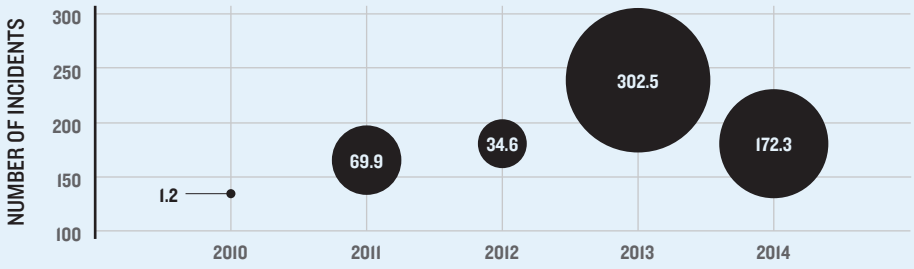

Method

Types of major industry security incidents from 2011 to 2013. Incidents may include theft of non-payment information.

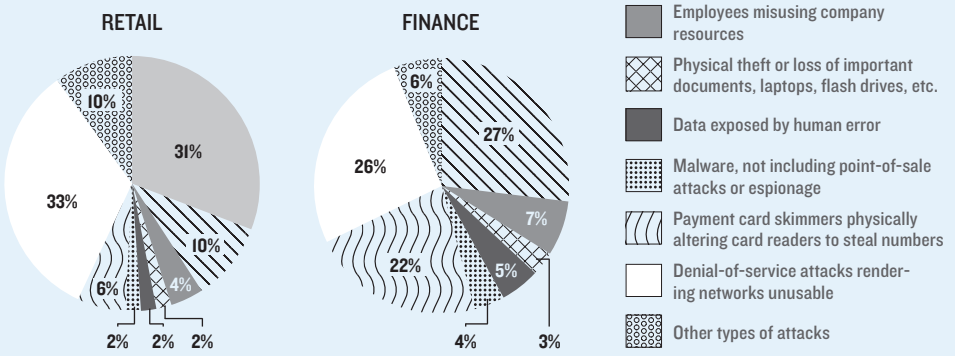

SOURCES: VERIZON 2014 DATA BREACH INVESTIGATIONS REPORT; RISK-BASED SECURITY / CYBER RISK ANALYTICS

Technology

A Weekend in Bitcoin City: Arnhem, the Netherlands

Is it possible to live for 48 hours on nothing but Bitcoin?

● Propped up next to me on the red sofa in my room at the boutique Hotel Modez in Arnhem, the Netherlands, my iPad has its screen on the Bitcoin exchange Bitstamp.net, and the value of the cryptocurrency is dropping, moment by moment. At breakfast one bitcoin had been worth over $400, but the value has been sinking for the past 30 minutes and has now hit $383. I know I'm blowing it.

When the price of a bitcoin tumbles by another $10, I am unwilling to risk any further losses. My stomach sinking, I head to reception to pay my bill. My exchange rate later turns out to have been near the day's low, and I feel like a sucker.

For the vast majority of Bitcoin holders—and the billions of people who have never even heard of the digital currency—such fluctuations may not seem like a big problem. But for me, sitting in the Hotel Modez, it was very real: I'd committed to paying for a room, priced in euros, with Bitcoin. As I waited, and the exchange worked against me, my bill had grown increasingly expensive.

Such is the state of affairs in the volatile world of cryptocurrencies—where regulation is a distant concept and large market swings are commonplace. Because it's hard to trace, Bitcoin has become common currency for criminals, but the list of legitimate companies accepting it as payment—or planning to—is growing to include the retailers Overstock and Newegg and the mainstream travel site Expedia, among others. Homes have been purchased with Bitcoin, which leans heavily on cryptography and a public ledger system called the blockchain. So has a hoped-for trip to space.

A rising number of people report or anticipate transacting in Bitcoin, and advocates see great potential in the currency for lowering the transaction cost of payments while increasing their security. But for Bitcoin to survive as a functional currency, it has to be widely accepted and useful in the way cash and credit cards are today. Retailers will need a reason to accept it—because of the lower costs, perhaps—and consumers will have to be convinced it's no more of a hassle than paying by conventional means. Can Bitcoin pass that test?

To find out, I had come to Arnhem, a place with one of the highest concentration of merchants accepting Bitcoin anywhere in the world. My experiment: Could a journalist plan a weekend escape paid for entirely with Bitcoin? Further, could he not only survive but perhaps even enjoy himself?

Arnhem's friendliness to Bitcoin has much to do with Patrick van der Meijde, a 36-year-old resident of this city of 150,000 located on the Rhine. Van der Meijde heard about Bitcoin a few years ago. Finding the concept intellectually interesting, and figuring that the traditional banking system could use competition, he decided to buy some. As his cache grew, he realized it wasn't so useful if he couldn't use it to buy things. So, with two partners, he set up a payment system local vendors could run—on their phones or any other connected device, like a laptop or tablet—allowing the owners to accept Bitcoin but be paid in euros. Van der Meijde has now convinced 45 businesses to accept Bitcoin, including a hotel and a major franchise grocery store.

Bitcoin has become common currency for criminals, but the list of legitimate companies accepting it for payment is growing.

Step one: Buying the plane ticket

Though I was familiar with Bitcoin, its genesis, its technical underpinnings, and its controversies, I didn't actually own any of the money. So eight days before I set out for Arnhem, I opened an account with a Boston-based startup called Circle that would let me buy bitcoins with a credit card.

Next, I logged onto CheapAir.com, one of a few companies that will let you book flights using Bitcoin, and bought a ticket flying KLM to Amsterdam from Munich. At the payment page I chose the option to display a Bitcoin address—a 25-to-34-character string of letters and numbers—to which I could send my payment. I then logged into Circle again to buy enough bitcoins to cover the ticket, but the transaction was immediately denied. After a call to my bank to explain that the charge indeed was not fraudulent, I tried again. This time I bought $450 worth of bitcoins, safely within Circle's $500-per-week credit card limit. The transaction went through instantly.

Proud of being a part of the future, I went to Circle's payments page, entered the Bitcoin address from CheapAir, and keyed in the $450 listed as the cost of my ticket. Almost instantly CheapAir's website updated—to tell me I'd sent the wrong amount. What?

Was this a scam? I had taken screenshots at a few phases of paying, so I did a quick post-mortem and realized I'd made a beginner's mistake: CheapAir's price was listed in dollars with the equivalent in Bitcoin, so I'd also entered my payment via Circle in dollars. It felt like the intuitive thing to do—but it was wrong. Whether because of the volatility of Bitcoin or the fact that there are multiple exchanges to price it and they rarely match up, the payment I had sent was roughly $1.60 short.

I started making calls. Charlie at Circle, stumped, suggested a do-over. Gemma at CheapAir was sure we could resolve the problem but insisted, "Only our CEO has access to the Bitcoin stuff."

A Bitcoin Weekend

Highlights of 48 hours living on cryptocurrency

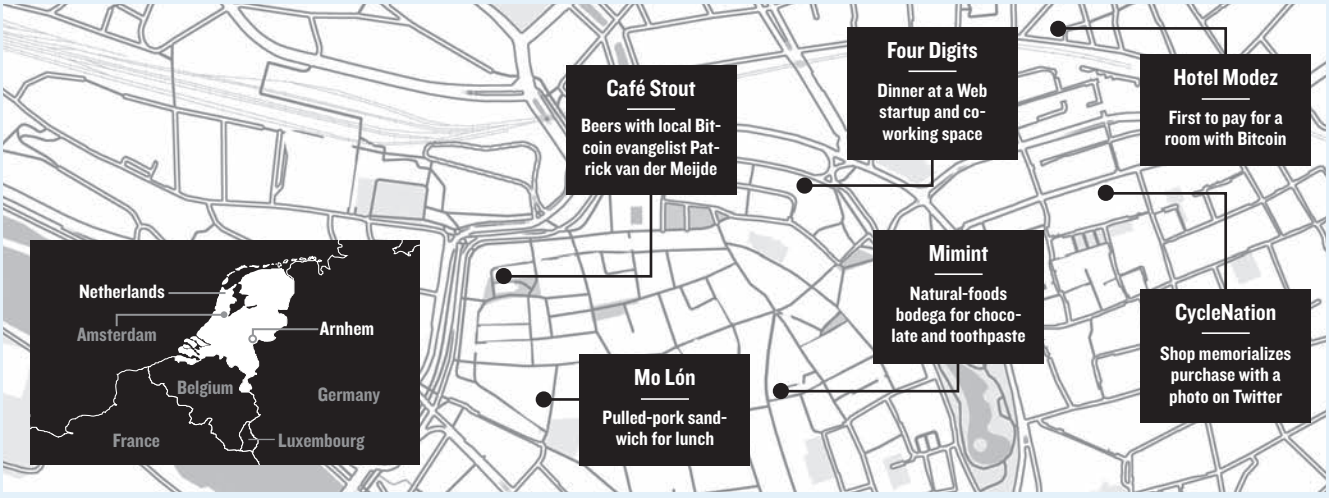

She told me to wait for CEO Jeff Klee to get in for the day—and he'd take care of it for me.

About an hour later, I received an e-mail confirmation for my flight. "Do I owe you five bucks or something?" I asked Gemma when I called her back. "Don't worry about it," she said. "It was easier to just issue the ticket."

Using Bitcoin for everyday purchases was proving more difficult than paying with a credit card.

Step two: Bitshock

A part of me expected Arnhem, about an hour from Amsterdam by train, to feel like a high-tech hub. But instead it resembled any typical European city. It had a few churches, a central pedestrian area filled with shops, and a handful of antique Dutch windmills. After checking into Hotel Modez, where the bubbly owner said I would be the first customer to pay in Bitcoin (something I'd hear a few times over the weekend; at the CycleNation bike shop a bemused employee took a photo of me at the register and posted it to Twitter), I met up with van der Meijde at a bar named Stout to chat over a beer.

"With Bitcoin?" said the gray-haired bartender when it came time to pay. He knew van der Meijde's face, as did others in town: he is variously referred to as "this guy who is really into Bitcoin" or, more simply, "that Bitcoin guy."

Payment itself was seamless: the bartender pulled up a QR code on his phone, van der Meijde scanned it using a Bitcoin wallet app called Mycelium on his phone, and the payment registered instantly. Later we repeated the process as I transferred Bitcoin directly to van der Meijde to cover my drinks.

A befuddled kid at the bar, of around university age, wanted to know what we were doing. "You mean I can buy drinks with this?" the kid asked. "Yes, of course," said van der Meijde. Ever the evangelist, he helped the kid download a Bitcoin wallet—and then transferred five euros' worth of bitcoins to him. The kid's friend watched all this with a stunned look on his face. "Does Café 'T Huys take this?" the kid asked. When van der Meijde told him that the bar did, the kid bolted for the door with his phone, now a few millibits richer, clutched in his hand.

Step three: Full Dutch crypto (mostly)

Over the next two days, guided by a map hosted on van der Meijde's website, I did indeed spend nothing but Bitcoin—with fewer mishaps than early users of Apple Pay were reporting in the press at about the same time. At dinner I ate a massive pile of ribs and learned that tips in Bitcoin are handled much like tips using a credit card, with waiters paid out from the register. Another day, at a restaurant called Mo Lón, I ordered a heaping pulled-pork sandwich and scanned a QR code that the owner, a Bitcoin enthusiast, loaded on a large LCD TV on the wall. At Mimint, a natural-foods bodega, I bought chocolate and toothpaste.

At Hotel Modez I was the first customer to pay in Bitcoin. At the CycleNation bike shop a bemused employee took a photo of me at the register and posted it to Twitter.

In only a few places did I encounter obstacles. At a souvenir shop I had to wait a few minutes for the owner to arrive, since he was the only one who knew how to accept Bitcoin. And at another shop I had some momentary Wi-Fi problems. I was turned away only once, at a small restaurant where the young woman work-

ing that day hadn't heard of Bitcoin. The cook, who was sitting at a table waiting for customers, said, "I've heard of Bitcoin, but I don't think we take it. Maybe the previous owners did?" (They were both surprised when I showed them the Bitcoin sticker affixed next to ones from MasterCard and Visa on the restaurant's window.)

One night, at van der Meijde's suggestion, I stopped by unannounced at a Web

45
Number of businesses in Arnhem that take bitcoins

startup and co-working space called Four Digits, where once a week around a dozen tech-minded folks in Arnhem get together informally to eat, drink, and geek out. A few of the people there were equal parts skeptical and excited by the idea of cryptocurrencies. Two of them discussed theoretical—but exceedingly unlikely—exploits that would let them rip off the point-of-sale application van der Meijde had helped design. The Indian delivery order they were eating had been paid for with Bitcoin, and as they split the bill, some were paying the purchaser back in Bitcoin, too.

"How much did dinner cost?" I asked.

"About half a bitcoin," someone said.

Step four: Ex post crypto

Though I mostly enjoyed the weekend, the getaway to Arnhem at times felt like a chore. I'd had to skip some cultural sites recommended by a friend from Amsterdam. The Hoge Veluwe National Park and its museum with works by van Gogh, Rodin, and Dubuffet don't accept Bitcoin, sadly.

Having exhausted most of the possible Bitcoin-ready diversions in town, I spent the last few hours of my visit, on a rainy Sunday, walking along the waterfront and through a park. I pined for a museum, or a bowling alley, or a film in a warm theater. Spending bitcoins had been easy, and ultimately—despite the snafu paying for the hotel—not that expensive. But the options had run out quickly.

And there was one exceedingly important thing I could not do: get out of town. The only way to travel between Arnhem and the airport using Bitcoin was to rent a car or hire a taxi—a multi-hundred-euro expense. By contrast, the train ticket to the same place, which was not payable by Bitcoin, cost only 17.10 euros. Even the very staunchest Bitcoin enthusiast would be unlikely to pay that kind of premium.

Knowing this, I had come to town with 18 euros in my pocket. I felt I was betraying something or someone—perhaps Bitcoin's mysterious inventor, Satoshi Nakamoto—as I inserted coin after real-euro coin in the yellow ticket machine of the government-owned Dutch Railways. When the railroad accepts Bitcoin, I thought, we'll know that cryptocurrencies have truly arrived. *—Russ Juskalian*

Q&A

Increasing the GDP of the Internet

Improving payment technology is vital if Internet companies are to make an impact, says Stripe CEO Patrick Collison.

● When Apple launched its new payment service last October, it boasted of support from major partners such as American Express, Chase, and Macy's. But Apple had also spent months ahead of the launch working with a relative minnow in the world of financial technology: a four-year-old San Francisco startup called Stripe.

Stripe got started in 2010 selling tools that make it easy for businesses to add credit card payment functionality to a website or mobile app. It quickly earned a reputation for being friendlier to coders than more traditional payment processors, making it popular with app makers and a natural choice for Apple when it began designing Apple Pay.

Today, Stripe's products have expanded to include subscription billing services and an online checkout system. Customers include Walmart, Twitter, and the ride-sharing app Lyft. The company has received $190 million from venture funders including PayPal cofounders Peter Thiel, Elon Musk, and Max Levchin.

In an interview with *MIT Technology Review*'s San Francisco bureau chief, Tom Simonite, in the cafeteria at his company's headquarters in the Mission district of San Francisco, CEO and cofounder Patrick Collison explains how he is trying to increase the small fraction of spending that takes place through e-commerce.

Payments processing seems like a fairly simple function. What makes it important?

Today the most exciting technology companies are the mobile marketplaces, companies like Airbnb and Lyft. Software is suffusing every industry and sector and market. As technology companies expand into more markets that have been traditionally offline, it's natural that business models are more about payments. Stripe is providing the infrastructure for the next wave of these companies.

Only 2 percent of commerce worldwide is e-commerce today. Why is that share so small?

There are major infrastructural deficiencies. If you're in Latin America or China and go to a website, it's almost guaranteed that you can't buy from it. By only accepting credit cards, which is what the majority of websites do, they're essentially restricting themselves to selling to North American and Western European buyers.

Even in the U.S., e-commerce only accounts for just over 6 percent of all retail. Is that a business problem or a technology problem?

It's absolutely a technology problem. Think of the places where you're likely to discover something that you want to purchase. You no longer stumble across it when you're walking down the high street; you find it in your Facebook or Twitter feed, on your phone.

But think about what you're supposed to do on a mobile device: click on the link, get bounced to some random e-commerce website, click "add to cart," zoom in, peck in your address and credit card number, and click on checkout. It's a 10-step process, and it might not even work at the end of it all. I can hand a dollar to somebody really easily in the store; it's really difficult for me to digitally hand them a dollar.

Apple Pay, which you worked on, seems to work well, and there are a lot of partners. But it's not very radical, is it?
I think Apple was quite aggressive on a number of points. Giving merchants a "token" versus your credit card number is a very important shift. If that business is compromised, you're not liable to anyone in the world stealing from you. The privacy changes are not insignificant. It would have been very natural for Apple to try to aggregate all this data and use it for ad targeting, but they haven't.

There's no point proposing some utopian vision that doesn't come to pass. A new, successful, behavior-altering payments product is a big deal and represents a crack in that frozen ice of the industry and the way things are done.

Stripe is invested in some more radical ideas. You support payments made via Bitcoin, and you've invested $3 million in a Bitcoin alternative called Stellar.
Bitcoin is kind of a financial Rorschach test; everyone projects their desired monetary future onto it. What we care about is making digital transactions and commerce more universal. There is the element to Bitcoin of just being a universal means to transport value.

Bitcoin has some user experience issues. Transactions take minutes to clear. There's the difficulty of obtaining bitcoins. With Stellar you can use real, normal currencies in addition to digital ones. Transactions clear instantly. We backed this because we're in favor of anything that seems it might help us build a more ubiquitous, useful commerce infrastructure for the Internet.

Can you really make as much money as other high-margin software startups by taking a small slice of transactions you handle?
The economics of the business generally work better than people think. Look at PayPal's income statement and margins—they're really good. We're a technology provider and happen to bill for that as a function of the transaction volume, but that's an implementation detail of the pricing. Do you price it as a percentage of the transaction or as a monthly fee or an annual fee? There's no reason why the margins should be that different.

"A new, successful, behavior-altering payments product is a big deal and represents a crack in that frozen ice of the industry and the way things are done."

—Patrick Collison, CEO of Stripe

Case Study

Starbucks Bets the Store on Mobile

The iconic coffee merchant built a payment app that quickly became the envy of retail. Can it recharge Starbucks's growth despite new competition?

● If pressed to name the leading smartphone app for paying for a purchase in a store, would you say Apple Pay? PayPal? Google Wallet?

Something from Visa, MasterCard, or one of the big banks, perhaps?

No doubt you've guessed it's none of the above. Starbucks has emerged as the runaway leader in mobile payments at retail stores. Twelve million active users now pay for their Frappuccino with a wave of a phone. Of the $1.6 billion spent via smartphone in U.S. stores in 2013, the company claims, a full 90 percent went to Starbucks, and most payments experts don't doubt it.

Starbucks's mobile wallet is actually a digital stored-value card more akin to its popular gift card. Its success—the app now accounts for 16 percent of Starbucks's 47 million weekly transactions (up 50 percent from a year ago)—makes it both a model and a target for payment apps from other retailers and tech companies alike.

Starbucks CEO Howard Schultz has even bigger plans for mobile payments. In comments to investors in late 2014, Schultz made it clear that he views the mobile app as a linchpin of plans to keep growth perking. For one thing, the company says, it speeds payment and reduces waits in line. It also makes it easier to track rewards points, giving customers incentive to buy more at Starbucks so they can get perks such as a free drink after buying 12.

But the real attraction is in the data that the company can collect. Whoever last touches the customer at the point of sale owns a wealth of information about purchases and preferences that can be fed back into the business. Starbucks has started using purchase data to send personalized offers the way merchants do online—say, for food to complement a drink. It's beginning to result in additional purchases. "It's going to be the Amazon experience, inside the store," says Ken Morris, a principal with Boston Retail Partners, which helps retailers create systems like Starbucks's.

To get more people to use the app and join the rewards program, the company plans to add major new features in 2015. One is the ability to order and pay in advance, which launched in Portland, Oregon, in December and will expand to other cities later. And in a few markets

starting in mid-2015, the app will enable customers to schedule regular delivery of individual drinks and food to offices—further expanding Starbucks's brand beyond its own stores.

Schultz wants to push deeper into the digital realm. He's talking to tech firms and other retailers about licensing its payment and rewards software, much the way Seattle neighbor Amazon sells its cloud computing services to other companies. Starbucks even aims to turn its stored-value cards, app, and rewards points into a broader currency usable at other retailers, creating a potential alternative to credit and debit cards.

Launched widely in January 2011, the app generates a bar code on the smartphone that can be scanned at checkout readers already installed in most stores. At the time, bar codes were pooh-poohed by tech firms already looking to a radio-based technology called near-field communication that Apple Pay now uses to send payment information. But Starbucks was looking for ease of use above all, says Chuck Davidson, who then worked for the company's prepaid-card business (worth $1.5 billion at the time) and is now head of customer engagement at the mobile-commerce firm CardFree.

The strategy worked. Just two months later, some three million people had paid with the app. Of course, Starbucks benefited from its brand name, customers who can afford smartphones, and a daily-habit product. But Davidson and other payment experts say other features were the real key to its success. Allowing easy tracking of card balances and rewards points encouraged customers to use the app more often.

New alternatives like Apple Pay and rival retailers' apps present challenges for Starbucks. Dunkin' Donuts, for instance, less than a year after relaunching its mobile app with a new rewards program in January 2014, said it reached two million rewards members and 10 million app downloads.

Starbucks created the iconic mobile payment app. Now its customers know how to use their phone to pay everywhere else, too. *—Robert D. Hof*

Case Study

Can Mobile Money Conquer New Markets?

After great success in East Africa, mobile-phone wallet M-Pesa hit some bumps in South Africa.

● Since its launch in 2007 by mobile-phone giant Vodafone and local partners, M-Pesa has taken East Africa by storm. Today more than 18 million users, most in Kenya and Tanzania, use the mobile-phone wallet to transfer billions of dollars a month. The technology has brought new financial options to places where banks and credit cards are scarce and cash has long been king, promising a faster, cheaper, more secure way to pay for things and send money.

The service enables cell-phone owners to essentially use their mobile device like a bank card. After registering as a user with an M-Pesa agent, a customer can upload money onto the phone. Those funds can then be used for many transactions, from grocery purchases to paying utility bills.

According to the Central Bank of Kenya, the value of M-Pesa transactions in Kenya jumped 30 percent, to $12 billion, in the first six months of 2014 compared with the same period in 2013. In Tanzania, a country half as populous, monthly M-Pesa transactions are valued at $820 million.

But despite this remarkable record, M-Pesa has found that moving the technology into a different market can be a challenge. Though the service is now available in 10 countries, more than 100,000 of M-Pesa's 186,000 authorized agents worldwide are still in Kenya.

For clues to how M-Pesa can find broader acceptance, many are watching its 2014 relaunch in South Africa, where an initial introduction in 2010 fell far below expectations, signing up 100,000 users instead of the 10 million anticipated.

The 2010 version was a "carbon clone" of its Kenyan counterpart, not well suited to South African customers, says Herman Singh, the managing executive for mobile commerce at Vodacom, who is responsible for the South Africa relaunch. (Vodacom is 65 percent owned by Vodafone.) The rollout also suffered because there weren't enough agents to help customers upload and download.

Now M-Pesa has simplified the registration process for agents and increased their numbers in South Africa from 800 to 8,000. Users can use their mobile wallet at all retailers. And there is now a voucher system similar to the prepaid mobile-phone credit system most South Africans use to pay for cell time.

Vodacom is betting that M-Pesa's low cost will win over consumers. There are no monthly account fees, and most services, like cash deposits and electronic transfers from bank accounts, are free of charge.

Early results are encouraging. In the first four months following the relaunch, the number of users grew from 100,000 to 650,000, and more transactions have been processed than in M-Pesa's first four years in the country combined, Singh says.

—Miriam Mannak

ACCESS THE FULL REPORT ONLINE

Additional stories offer a deeper look at payment technology, cybercurrency, and the challenges faced by digital money and payments:

technologyreview.com/business

Reviews

Our Fear of Artificial Intelligence

Many technologists insist we need to be working now to ensure that computers smarter than us don't ruin the world. Are they making erroneous assumptions about what machines could someday do?

By Paul Ford

Years ago I had coffee with a friend who ran a startup. He had just turned 40. His father was ill, his back was sore, and he found himself overwhelmed by life. "Don't laugh at me," he said, "but I was counting on the singularity."

My friend worked in technology; he'd seen the changes that faster microprocessors and networks had wrought. It wasn't that much of a step for him to believe that before he was beset by middle age, the intelligence of machines would exceed that

Illustration by Jacob Escobedo

of humans—a moment that futurists call the singularity. A benevolent superintelligence might analyze the human genetic code at great speed and unlock the secret to eternal youth. At the very least, it might know how to fix your back.

But what if it wasn't so benevolent? Nick Bostrom, a philosopher who directs the Future of Humanity Institute at the University of Oxford, describes the following scenario in his book *Superintelligence*, which has prompted a great deal of debate about the future of artificial intelligence. Imagine a machine that we might call a "paper-clip maximizer"—that is, a machine programmed to make as many paper clips as possible. Now imagine that this machine somehow became incredibly intelligent. Given its goals, it might then decide to create new, more efficient paper-clip-manufacturing machines—until, King Midas style, it had converted essentially everything to paper clips.

Superintelligence: Paths, Dangers, Strategies
By Nick Bostrom
Oxford University Press, 2014

No worries, you might say: you could just program it to make exactly a million paper clips and halt. But what if it makes the paper clips and then decides to check its work? Has it counted correctly? It needs to become smarter to be sure. The superintelligent machine manufactures some as-yet-uninvented raw-computing material (call it "computronium") and uses that to check each doubt. But each new doubt yields further digital doubts, and so on, until the entire earth is converted to computronium. Except for the million paper clips.

Bostrom does not believe that the paper-clip maximizer will come to be, exactly; it's a thought experiment, one designed to show how even careful system design can fail to restrain extreme machine intelligence. But he does believe that superintelligence could emerge, and while it could be great, he thinks it could also decide it doesn't need humans around. Or do any number of other things that destroy the world. The title of chapter 8 is: "Is the default outcome doom?"

If this sounds absurd to you, you're not alone. Critics such as the robotics pioneer Rodney Brooks say that people who fear a runaway AI misunderstand what computers are doing when we say they're thinking or getting smart. From this perspective, the putative superintelligence Bostrom describes is far in the future and perhaps impossible.

Yet a lot of smart, thoughtful people agree with Bostrom and are worried now. Why?

Volition

The question "Can a machine think?" has shadowed computer science from its beginnings. Alan Turing proposed in 1950 that a machine could be taught like a child; John McCarthy, inventor of the programming language LISP, coined the term "artificial intelligence" in 1955. As AI researchers in the 1960s and 1970s began to use computers to recognize images, translate between languages, and understand instructions in normal language and not just code, the idea that computers would eventually develop the ability to speak and think—and thus to do evil—bubbled into mainstream culture. Even beyond the oft-referenced HAL from *2001: A Space Odyssey*, the 1970 movie *Colossus: The Forbin Project* featured a large blinking mainframe computer that brings the world to the brink of nuclear destruction; a similar theme was explored 13 years later in *WarGames*. The androids of 1973's *Westworld* went crazy and started killing.

When AI research fell far short of its lofty goals, funding dried up to a trickle, beginning long "AI winters." Even so, the torch of the intelligent machine was carried forth in the 1980s and '90s by sci-fi authors like Vernor Vinge, who popularized the concept of the singularity; researchers like the roboticist Hans Moravec, an expert in computer vision; and the engineer/entrepreneur Ray Kurzweil, author of the 1999 book *The Age of Spiritual Machines*. Whereas Turing had posited a humanlike intelligence, Vinge, Moravec, and Kurzweil were thinking bigger: when a computer became capable of independently devising ways to achieve goals, it would very likely be capable of introspection—and thus able to modify its software and make itself more intelligent. In short order, such a computer would be able to design its own hardware.

As Kurzweil described it, this would begin a beautiful new era. Such machines would have the insight and patience (measured in picoseconds) to solve the outstanding problems of nanotechnology and spaceflight; they would improve the human condition and let us upload our consciousness into an immortal digital form. Intelligence would spread throughout the cosmos.

Extreme AI predictions are "comparable to seeing more efficient internal combustion engines ... and jumping to the conclusion that warp drives are just around the corner," Rodney Brooks writes.

You can also find the exact opposite of such sunny optimism. Stephen Hawking has warned that because people would be unable to compete with an advanced AI, it "could spell the end of the human race." Upon reading *Superintelligence*, the entrepreneur Elon Musk tweeted: "Hope we're not just the biological boot loader for digital superintelligence. Unfortunately, that is increasingly probable." Musk then followed with a $10 million grant to the Future of Life Institute. Not to be confused with Bostrom's center, this is an organization that says it is "working to mitigate existential risks facing humanity," the ones that could arise "from the development of human-level artificial intelligence."

No one is suggesting that anything like superintelligence exists now. In fact, we still have nothing approaching a general-purpose artificial intelligence or even a clear path to how it could be achieved. Recent advances in AI, from automated assistants such as Apple's Siri to Google's driverless cars, also reveal the technology's severe limitations; both can be thrown off by situations that they haven't encountered before. Artificial neural networks can learn for themselves to recognize cats in photos. But they must be shown hundreds of thousands of examples and still end up much less accurate at spotting cats than a child.

This is where skeptics such as Brooks, a founder of iRobot and Rethink Robotics, come in. Even if it's impressive—relative to what earlier computers could manage—for a computer to recognize a picture of a cat, the machine has no volition, no sense of what cat-ness is or what else is happening in the picture, and none of the countless other insights that humans have. In this view, AI could possibly lead to intelligent machines, but it would take much more work than people like Bostrom imagine. And even if it

Events

SXSW
March 13-22, 2015
Austin, TX
http://sxsw.com/

EmTech China
Beijing, China
www.technologyreview.com/events

EmTech Digital
June 1-2, 2015
San Francisco, CA
www.technologyreview.com/emtech/digital/15/

BIO International Convention
June 15-18, 2015
Philadelphia, PA
http://convention.bio.org/

EmTech MIT
October 6-8, 2015
Cambridge, MA
www.emtechmit.com

could happen, intelligence will not necessarily lead to sentience. Extrapolating from the state of AI today to suggest that superintelligence is looming is "comparable to seeing more efficient internal combustion engines appearing and jumping to the conclusion that warp drives are just around the corner," Brooks wrote recently on Edge.org. "Malevolent AI" is nothing to worry about, he says, for a few hundred years at least.

Insurance policy

Even if the odds of a superintelligence arising are very long, perhaps it's irresponsible to take the chance. One person who shares Bostrom's concerns is Stuart J. Russell, a professor of computer science at the University of California, Berkeley. Russell is the author, with Peter Norvig (a peer of Kurzweil's at Google), of *Artificial Intelligence: A Modern Approach*, which has been the standard AI textbook for two decades.

"There are a lot of supposedly smart public intellectuals who just haven't a clue," Russell told me. He pointed out that AI has advanced tremendously in the last decade, and that while the public might understand progress in terms of Moore's Law (faster computers are doing more), in fact recent AI work has been fundamental, with techniques like deep learning laying the groundwork for computers that can automatically increase their understanding of the world around them.

Bostrom's book proposes ways to align computers with human needs. We're basically telling a god how we'd like to be treated.

Because Google, Facebook, and other companies are actively looking to create an intelligent, "learning" machine, he reasons, "I would say that one of the things we ought not to do is to press full steam ahead on building superintelligence without giving thought to the potential risks. It just seems a bit daft." Russell made an analogy: "It's like fusion research. If you ask a fusion researcher what they do, they say they work on containment. If you want unlimited energy you'd better contain the fusion reaction." Similarly, he says, if you want unlimited

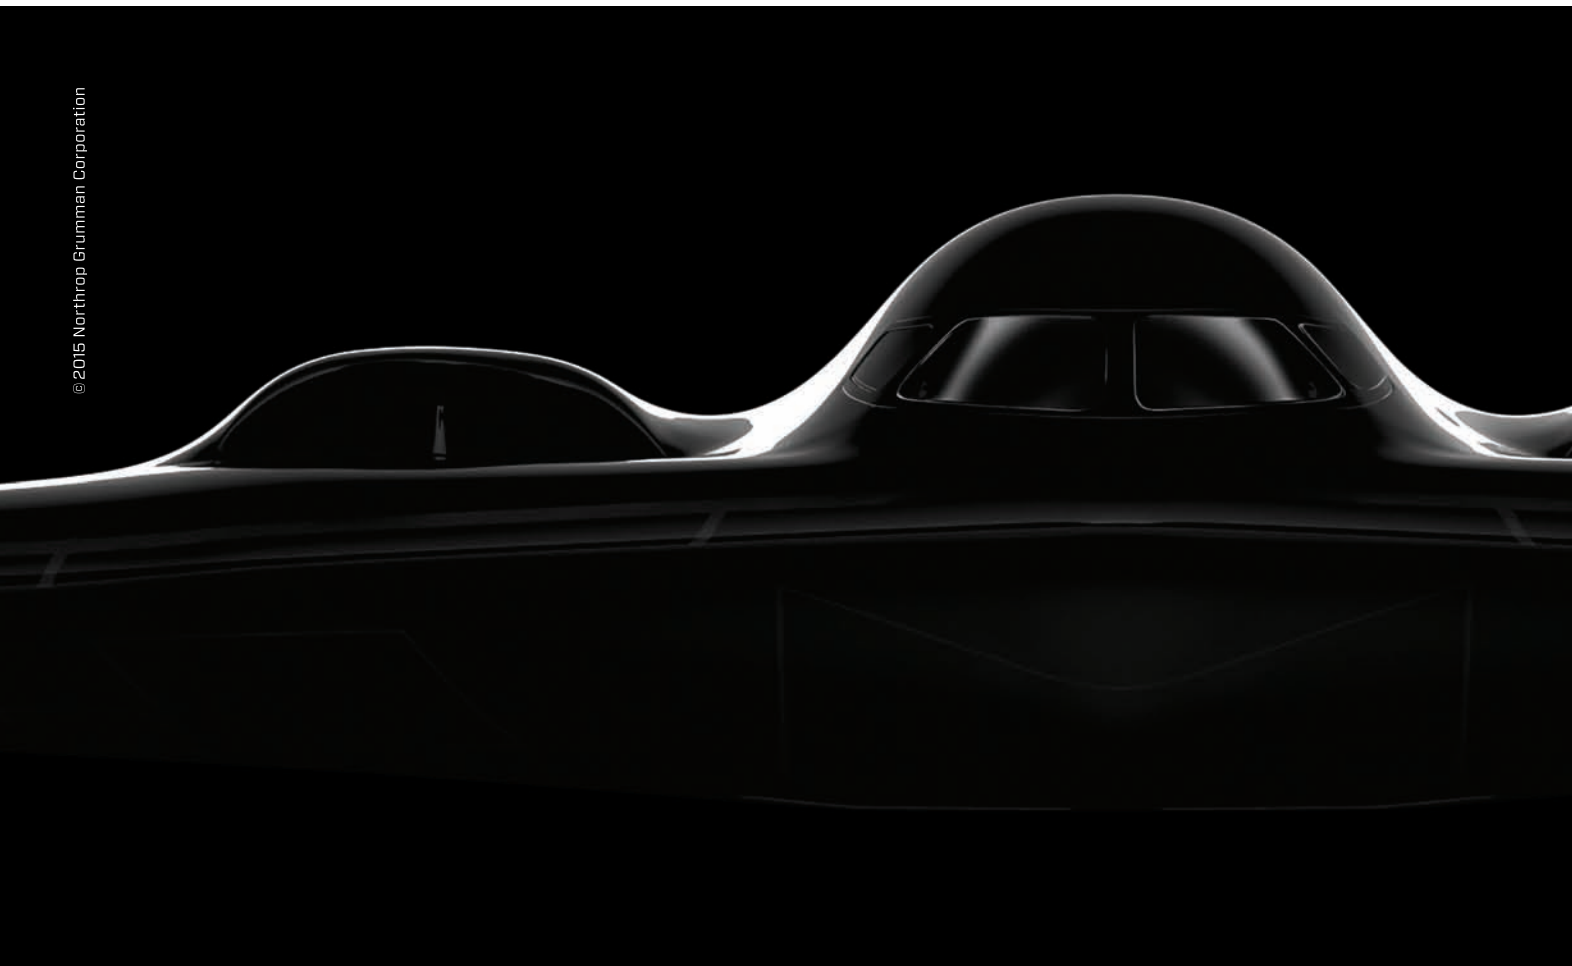

intelligence, you'd better figure out how to align computers with human needs.

Bostrom's book is a research proposal for doing so. A superintelligence would be godlike, but would it be animated by wrath or by love? It's up to us (that is, the engineers). Like any parent, we must give our child a set of values. And not just any values, but those that are in the best interest of humanity. We're basically telling a god how we'd like to be treated. How to proceed?

Bostrom draws heavily on an idea from a thinker named Eliezer Yudkowsky, who talks about "coherent extrapolated volition"—the consensus-derived "best self" of all people. AI would, we hope, wish to give us rich, happy, fulfilling lives: fix our sore backs and show us how to get to Mars. And since humans will never fully agree on anything, we'll sometimes need it to decide for us—to make the best decisions for humanity as a whole. How, then, do we program those values into our (potential) superintelligences? What sort of mathematics can define them? These are the problems, Bostrom believes, that researchers should be solving now. Bostrom says it is "the essential task of our age."

For the civilian, there's no reason to lose sleep over scary robots. We have no technology that is remotely close to superintelligence. Then again, many of the largest corporations in the world are deeply invested in making their computers more intelligent; a true AI would give any one of these companies an unbelievable advantage. They also should be attuned to its potential downsides and figuring out how to avoid them. This somewhat more nuanced suggestion—without any claims of a looming AI-mageddon—is the basis of an open letter on the website of the Future of Life Institute, the group that got Musk's donation. Rather than warning of existential disaster, the letter calls for more research into reaping the benefits of AI "while avoiding potential pitfalls."

This letter is signed not just by AI outsiders such as Hawking, Musk, and Bostrom but also by prominent computer scientists (including Demis Hassabis, a top AI researcher who is profiled on page 16). You can see where they're coming from. After all, if they develop an artificial intelligence that doesn't share the best human values, it will mean they weren't smart enough to control their own creations.

Paul Ford, a freelance writer in New York, wrote about Bitcoin in March/April 2014.

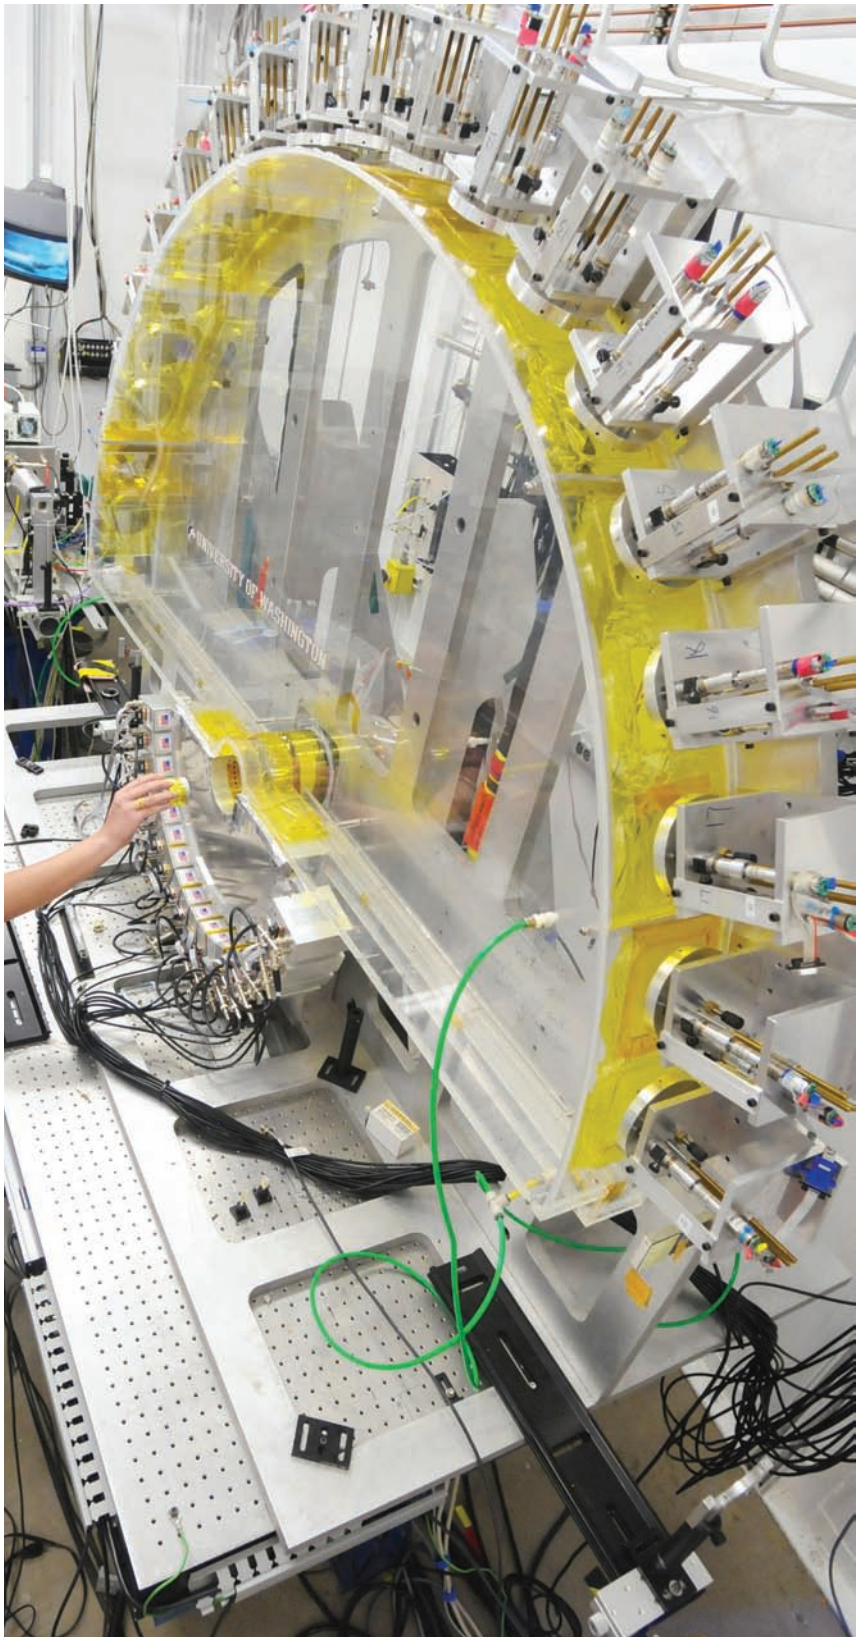

A lithium-ion battery is loaded into an x-ray device used to help researchers understand how experimental batteries work.

COURTESY OF ARGONNE NATIONAL LABORATORY

Why We Don't Have Battery Breakthroughs

A promising advance that came to nothing suggests what it will take to make cheap batteries for electric cars.

By Kevin Bullis

Electric cars are quick and quiet, with a range more than long enough for most commutes. If you want a car with extremely fast acceleration, the Tesla Model S is hard to beat. And, of course, electric vehicles avoid the pollution associated with conventional cars, including emissions of carbon dioxide from burning gasoline. Yet they account for a tiny fraction of automotive sales, mainly because the batteries that propel them are expensive and need to be recharged frequently.

A better battery could change everything. But while countless breakthroughs have been announced over the last decade, time and again these advances have failed to translate into commercial batteries with anything like the promised improvements in cost and energy storage. Some well-funded startups, most notably A123 Systems, began with bold claims but failed to deliver.

The Powerhouse
By Steve LeVine
Viking, 2015

The Powerhouse, a new book by journalist Steve LeVine, chronicles the story behind one of the most dramatic battery announcements of recent years and explains how it came to nothing. The announcement was made in February 2012, at a conference in Washington, D.C., where a crowd of researchers, entrepreneurs, and investors had come to hear the likes of Bill Gates and Bill Clinton expound on the importance of new energy technology—and also to tap into one of the newest funding sources in Washington, the Advanced Research Projects Agency for Energy, or ARPA-E. Founded in 2009, ARPA-E had been tasked with identifying potentially transformational research. The head of that agency, Arun Majumdar, was ready to unveil one of its first major successes: a battery cell, developed by the startup Envia, that could store twice as much energy as a conventional one. The cost of a battery that could take a car from Washington to New York without recharging, Majumdar said, would fall from $30,000 to $15,000. Electric cars would become far more affordable and practical.

Within months, GM licensed the technology and signed an agreement to support its development, gaining the right to use any resulting batteries. The deal was potentially worth hundreds of millions of dollars to Envia, LeVine writes. But soon Envia was getting frustrated messages from GM engineers who couldn't reproduce the startup's results. The year after the announcement, the deal was scuttled. Envia's impressive battery had been a fluke.

LeVine's account of Envia's work shows why major progress in batteries is so hard to achieve and why startups that promise world-changing breakthroughs have struggled. Over the last decade we've seen remarkable improvements in this industry, but they've come largely from established companies steadily making small advances.

Envia's cell was a new type of lithium-ion battery. Invented in the late 1970s and early 1980s and commercialized in the 1990s, these batteries generate electrical current when lithium ions shuttle between two electrodes. Light but powerful, they have transformed portable electronics. Their use in electric cars, however, is recent. In the 1990s, GM used cheaper lead-acid batteries for its electric EV-1; each battery weighed a bulky 600 kilograms and delivered only 55 to 95 miles before it needed to be recharged. When Tesla Motors introduced one of the first lithium-ion-powered electric cars in 2008, it could go 250 miles on a charge, roughly three times farther than the EV-1. But the vehicle cost over $100,000, in large part because the batteries were so expensive. To cut costs, the lithium-ion-powered electric cars made today by companies such as Nissan and GM use small battery packs with a range of less than 100 miles.

While countless breakthroughs have been announced over the last decade, time and again these advances failed to translate into commercial batteries.

One difficult thing about developing better batteries is that the technology is still poorly understood. Changing one part of a battery—say, by introducing a new electrode—can produce unforeseen problems, some of which can't be detected without years of testing. To achieve the kinds of advances venture capitalists and ARPA-E look for, Envia incorporated not just one but two experimental electrode materials.

LeVine describes what went wrong. In 2006 Envia had licensed a promising material developed by researchers at Argonne National Laboratory. Subsequently, a major problem was discovered. The problem—which one battery company executive called a "doom factor"—was that over time, the voltage at which the battery operated changed in ways that made it unusable. Argonne researchers investigated the problem and found no ready answer. They didn't understand the basic chemistry and physics of the material well enough to grasp precisely what was going wrong, let alone fix it, LeVine writes.

With its experimental material for the opposite electrode, this one based on silicon, Envia faced another challenge. Researchers had seemingly solved the major problem with silicon electrodes—their tendency to fall apart. But the solution required impractical manufacturing techniques.

When Envia made its announcement in 2012, it seemed to have figured out how to make both these experimental materials work. It developed a version of the silicon electrode that could be manufactured more cheaply. And through trial and error it had stumbled upon a combination of coatings that stabilized the voltage of the Argonne material. Envia cofounder Sujeet Kumar "understood that the answer was a composite of coatings," LeVine writes. "But he still didn't know what the composite was arresting or why it succeeded in doing so." Since Envia was a startup with limited funds, he "didn't have

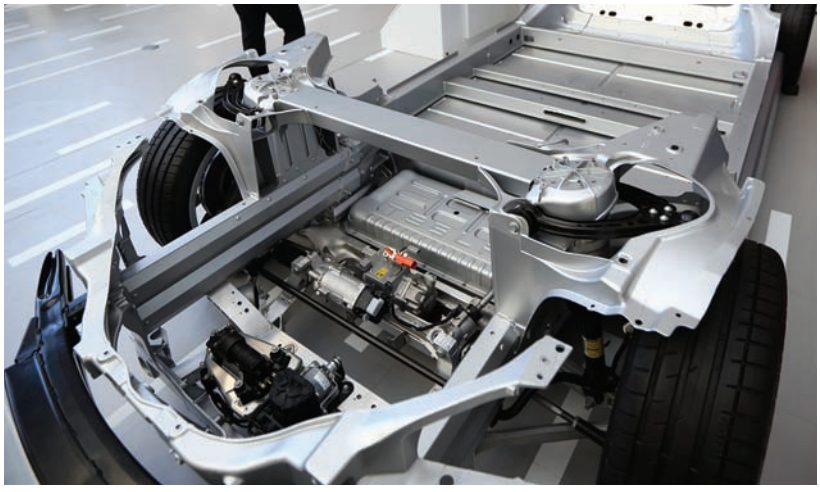

A look inside a Tesla Model S shows the battery, a gray slab that takes up most of the space between the front and rear wheels.

the instruments that could figure it out." But once it became obvious that the results Envia had reported for its battery couldn't be reproduced, understanding the problem became crucial. Even tiny changes to the composition of a material can have a significant impact on performance, so for all Envia knew, its record-setting battery worked because of a contaminant in a batch of material from one of its suppliers.

The story of Envia stands in sharp contrast to what's turned out to be the most successful recent effort to cut the price of batteries and improve their performance. This success hasn't come from a breakthrough but from the close partnership between Tesla Motors and the major battery cell supplier Panasonic. Since 2008, the cost of Tesla's battery packs has been cut approximately in half, while the storage capacity has increased by about 60 percent. Tesla didn't attempt to radically change the chemistry or materials in lithium-ion batteries; rather, it made incremental engineering and manufacturing improvements. It also worked closely with Panasonic to tweak the chemistry of existing battery materials according to the precise needs of its cars.

Tesla claims that it is on track to produce a $35,000 electric car with a roughly 200-mile range by 2017—a feat that's equivalent to what GM hoped to achieve with Envia's new battery. The company anticipates selling hundreds of thousands of these electric cars a year, which would be a big leap from the tens of thousands it sells now. Yet for electric cars to account for a significant portion of the roughly 60 million cars sold each year around the world, batteries will probably need to get considerably better. After all, 200 miles is far short of the 350-plus miles people are used to driving on a tank of gasoline, and $35,000 is still quite a bit more than the $15,000 price of many small gas-powered cars.

Since 2008, the cost of Tesla's battery packs has been cut approximately in half, while the storage capacity has increased by about 60 percent.

How will we close the gap? There is probably still plenty of room to improve lithium-ion batteries, though it's hard to imagine that Tesla's success with minor changes to battery chemistry will continue indefinitely. At some point, radical changes such as the ones Envia envisioned may be needed. But the lesson from the Envia fiasco is that such changes must be closely integrated with manufacturing and engineering expertise.

That approach is already yielding promising results with the Argonne material that Envia licensed. Envia's battery operated at high voltages to achieve high levels of energy storage. Now battery manufacturers are finding that using more modest voltage levels can significantly increase energy storage without the problems that troubled Envia. Meanwhile, battery researchers are publishing papers that show how trace amounts of additives change the behavior of the materials, making it possible to edge up the voltage and energy storage. The key is to combine research that illuminates details about the chemistry and physics of batteries with the expertise that battery manufacturers have gained in making practical products.

It's an industry in which it's very difficult for a startup, however enticing its technology, to go it alone. Andy Chu, a former executive at A123 Systems, which went bankrupt in 2012, recently told me why large companies dominate the battery industry. "Energy storage is a game played by big players because there are so many things that can go wrong in a battery," he said. "I hope startups are successful. But you can look at the history over the past few years, and it's not been good."

Kevin Bullis is MIT Technology Review*'s senior editor for materials.*

TOMOHIRO OHSUMI/BLOOMBERG VIA GETTY IMAGES

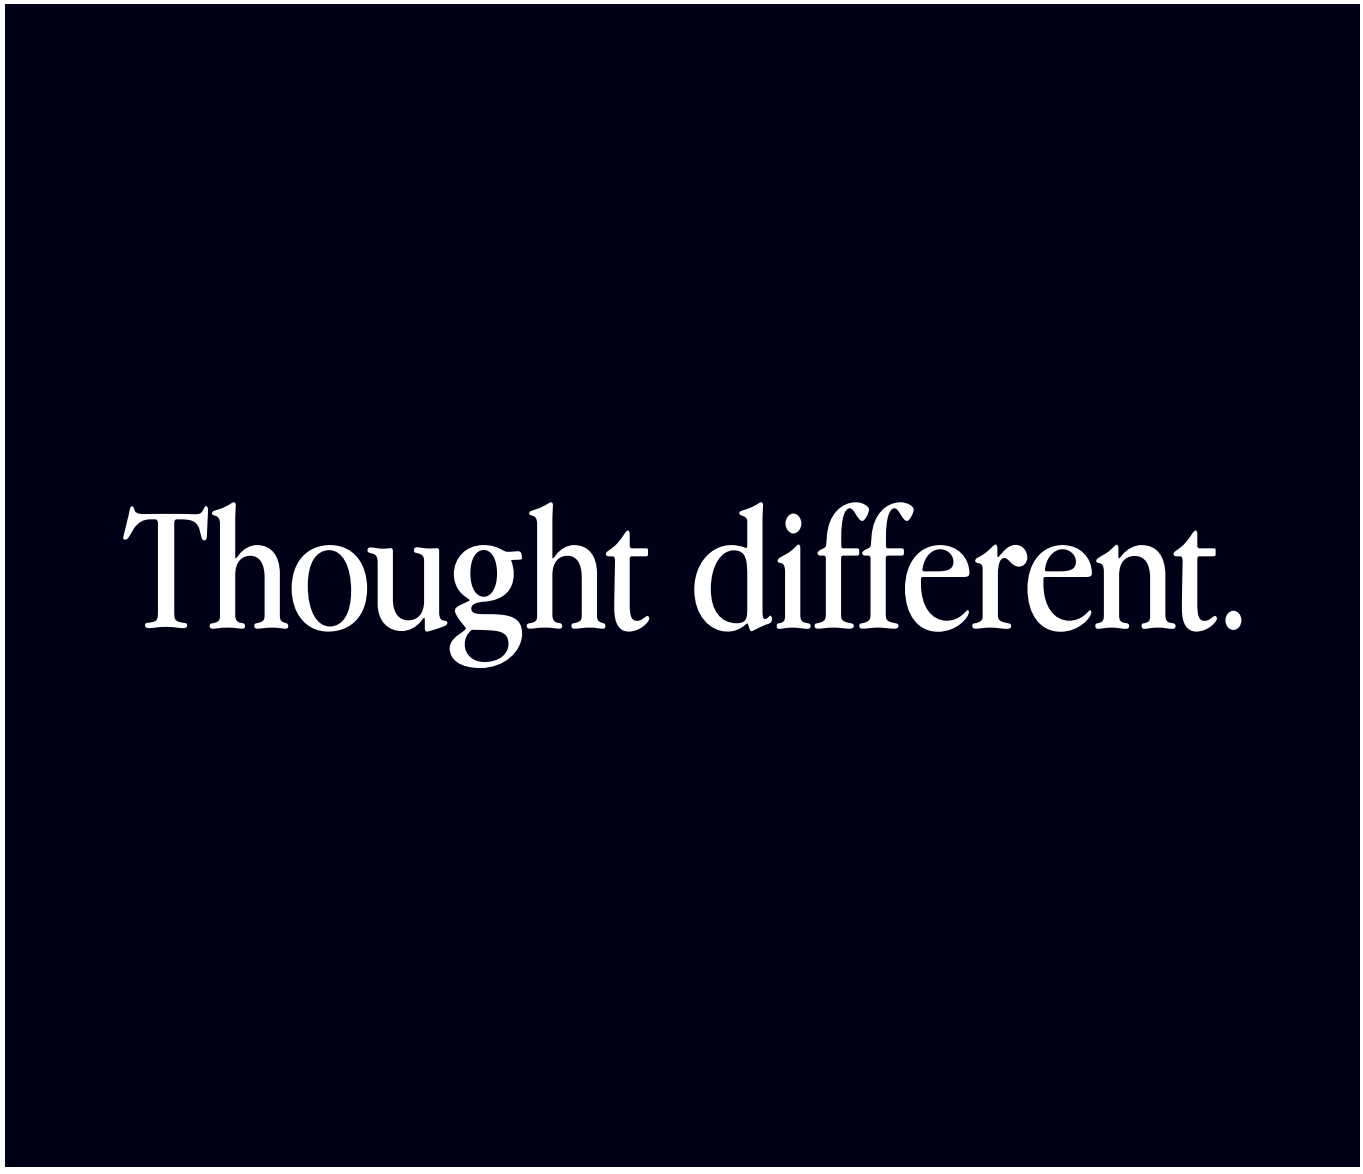

The Purpose of Silicon Valley

Capital and engineering talent have been flocking to seemingly trivial mobile apps. But would we really be better off if more startups instead went directly after big problems?

By Michael S. Malone

OTHER MEANS

The view from Mike Steep's office on Palo Alto's Coyote Hill is one of the greatest in Silicon Valley.

Beyond the black and rosewood office furniture, the two large computer monitors, and three Indonesian artifacts to ward off evil spirits, Steep looks out onto a panorama stretching from Redwood City to Santa Clara. This is the historic Silicon Valley, the birth-

place of Hewlett-Packard and Fairchild Semiconductor, Intel and Atari, Netscape and Google. This is the home of innovations that have shaped the modern world. So is Steep's employer: Xerox's Palo Alto Research Center, or PARC, where personal computing and key computer-networking technologies were invented, and where he is senior vice president of global business operations.

"This town used to think big—the integrated circuit, personal computers, the Internet. Are we really leveraging all that intellectual power and creativity creating Instagram and dating apps? Is this truly going to change the world?"

And yet Mike Steep is disappointed at what he sees out the windows.

"I see a community that acts like it knows where it's going, but that seems to have its head in the sand," he says. He gestures toward the Hewlett-Packard headquarters a few blocks away and Hoover Tower at Stanford University. "This town used to think big—the integrated circuit, personal computers, the Internet. Are we really leveraging all that intellectual power and creativity creating Instagram and dating apps? Is this truly going to change the world?"

After spending years at Microsoft, HP, and Apple, Steep joined PARC in 2013 to help the legendary ideas factory better capitalize on its work. As part of the job, he travels around the world visiting R&D executives in dozens of big companies, and increasingly he worries that the Valley will become irrelevant to them. Steep is one of 22 tech executives on a board the mayor of London set up to promote a "smart city"; they advise officials on how to allocate hundreds of millions of pounds for projects that would combine physical infrastructure such as new high-speed rail with sensors, databases, and analytics. "I know for a fact that China and an array of other countries are chasing this project, which will be the template for scores of similar big-city infrastructure projects around the world in years to come," Steep says. "From the U.S.? IBM. From Silicon Valley? Many in England ask if anyone here has even heard of the London subway project. That's unbelievable. Why don't we leverage opportunities like this here in the Valley?"

Steep isn't alone in asking whether Silicon Valley is devoting far too many resources to easy opportunities in mobile apps and social media at the expense of attacking bigger problems in energy, medicine, and transportation (see Q&A: Peter Thiel, November/December 2014). But if you put that argument to many investors and technologists here, you get a reasonable comeback: has Silicon Valley really ever set out to directly address big problems? In fact, the classic Valley approach has been to size up which technologies it can quickly and ambitiously advance, and then let the world make of them what it will. That is how we got Facebook and Google, and it's why the Valley's clean-tech affair was a short-lived mismatch. And as many people point out with classic Silicon Valley confidence, the kind of work that made the area great is still going on in abundance.

The next wave

A small group of executives, surrounded by hundreds of bottles of wine, sit in the private dining room at Bella Vita, an Italian restaurant in Los Altos's picturesque downtown of expensive tiny shops. Within a few miles, one can find the site of the original Fairchild Semiconductor, Steve Jobs's house, and the saloon where Nolan Bushnell set up the first Atari game. The host of this gathering is Carl Guardino, CEO of the Silicon Valley Leadership Group, an industry association dedicated to the economic health of the Valley. The 400 organizations that belong to the group are mostly companies that were founded long before the mobile-app craze; only 10 percent are startups. That is evident at this dinner, to which Guardino has invited three of his board members: Steve Berglund, CEO of Trimble, a maker of GPS equipment; Tom Werner, CEO of the solar provider SunPower; and Greg Becker, CEO of Silicon Valley Bank.

These are people who, like Steep, spend much of their time meeting with people in governments and other companies. Asked whether the Valley is falling out of touch with what the world really needs, each disagrees, vehemently. They are almost surprised by the question. "This is the most adaptive and flexible business community on the planet," says Becker. "It is always about innovation—and going where the opportunity leads next. If you're worried that the Valley is overpursuing one market or another, then just wait a while and it will change direction again. That's what we are all about."

"This is the center of world capitalism, and capitalism is always in flux," Werner adds. "Are there too many social-networking and app companies out there right now? Probably. But what makes you think it's going to stay that way for long? We have always undergone corrections. It's the nature of who we are … But we'll come out stronger than ever, and in a whole different set of markets and new technologies. This will still be the best place on the planet for innovation."

Berglund contends that a generational change already under way will reduce the emphasis on apps. "Young people don't seem to care as much about code as their generational elders," he says. "They want to build things—stuff like robots and drones. Just go to the Maker Faire and

watch them. They're going to take this valley in a whole different direction."

Berglund could be right. In the first half of 2014, according to CB Insights, Internet startups were the leading recipient of venture investment in San Francisco and Silicon Valley (the area got half of the U.S. total; New York was second at 10 percent). But investment in the Internet sector accounted for 59 percent of the total, down from a peak of 68 percent in 2011.

Doug Henton, who heads the consulting firm Collaborative Economics and oversaw an upcoming research report on the state of the Valley, argues that since 1950 the area has experienced five technological waves. Each has lasted about 10 to 20 years and encompassed a frenzy followed by a crash and shakeout and then a mature "deployment period." Henton has identified these waves as defense (1950s and 1960s), integrated circuits (1960s and 1970s), personal computers (1970s and 1980s), Internet (1990s), and social media (2000s and 2010s). By these lights, the social-media wave, however dominant it is in the public eye, soon may be replaced by another wave. Henton suggests that it's likely to involve the combination of software, hardware, and sensors in wearable devices and the "Internet of things."

Low-hanging fruit

Floyd Kvamme began his career at Fairchild Semiconductor (he was in a meeting with Gordon Moore and Andy Grove when an executive burst in to announce that John F. Kennedy had been assassinated), led the mainframe computer program at National Semiconductor, and then became a successful venture capitalist at Kleiner Perkins Caufield & Byers. Before he retired in 2009 to become an angel investor, he spent eight years as cochair of the President's Council of Advisors on Science and Technology.

To him, whether Silicon Valley should solve bigger problems is a misguided

question because it's already solving them indirectly. "The world's attention is always distracted by the newest big companies in the Valley," says Kvamme, "but the reality is one of ubiquity. The fact is that the Valley is into everything these days. Why? Because even after all these years, it still comes down to the silicon—and not just processors or memory but sensors, and most of all, because everything is mobile now, power integration and management.

"I don't think people appreciate the miracle that's taken place over the last decade. In 2005 [the Council of Advisors] did an energy report that determined that the U.S. was consuming 100 quadrillion BTUs of energy each year. We forecast that by 2030, the demand would grow to 150 quads. Instead, [in 2013] the U.S. only consumed 98 quads. How was this accomplished? By improved power control—and the low-hanging fruit was efficiency at the silicon and application level, Silicon Valley's strengths. Almost without the world noticing, efficient processors are showing up everywhere."

A lot of work in the Valley "seems trivial, small ball," but "the myth is that we are supposed to be out changing the world. In fact, our utopian idealism has shrunk in recent years as we've come to realize that even little innovations, in the right context, can have enormous impact."

Such combinations of improvements in hardware and software lead to many technologies that are extremely valuable even if they get less attention than Web or mobile-app startups valued at hundreds of millions of dollars, says Tom Hayes, a marketing executive who founded the Techmanity conference and Joint Venture Silicon Valley, a group that promotes regional development. What confuses outsiders is that so much in the Valley "seems trivial, small ball," Hayes says, "when the myth is that we are supposed to be out changing the world. In fact, our utopian idealism has shrunk in recent years as we've come to realize that even little innovations, in the right context, can have enormous impact … and the odds of pulling them off successfully are infinitely greater."

As an example of the "right context," Hayes cites cars. Autonomous vehicles, he believes, will change the nature of cities and become a new platform, as PCs and smartphones were. "This achievement alone will be enough to create another golden age in the Valley," he says. "And that's just one such revolution: drones will be another huge platform, as will mobile medical monitoring devices, including smart watches. And I have no doubt there will be more."

Linked together

Cofounding LinkedIn made Reid Hoffman a billionaire, but it is as a venture capitalist at Greylock Partners over the last five years that he has become a guru. On any given day, Greylock's lobby is jammed with as many as a dozen startup teams waiting to meet with Hoffman, who races from one conference room to the next.

Hoffman agrees with Steep's assessment that Valley technologists are fixated on quick opportunities in apps and software, but he sees that as appropriate, because software affects every level of products and organizations. That means it offers opportunities to have huge impacts.

"It's not accidental that Tesla is here, because the fundamental thing about its product isn't the battery but the software," he says. "So as software tackles one industry after another at the structural level, the companies created will be mostly based here. Similarly, when software defines human ecosystems, especially in English, they too will largely be based here in the Valley. Just look at Twitter, Facebook, eBay, Dropbox, and all of the others. It's not a coincidence that they are in the Valley. Aggregators too, such as Pinterest. Big data will be located in many places, but most of the big-data analytical toolmakers will be here as well. So will the next generation of [operating-system] layers—just as Android and [Apple's] iOS are now."

This means that there will always be vast swaths of technology where Valley companies don't compete and big problems for people elsewhere to try to solve. Those challenges don't mesh with Valley VCs' expectations that companies start quickly, with low capital costs, and have the potential to scale up tremendously. "There's a lot of things that we aren't good at here," Hoffman says, "and to which we have an almost antibody weakness. We don't know how to cultivate a DuPont-type company here. Health? Genetics? Probably they will remain centered in Boston, unless [those industries] develop a strong entrepreneurship component. Then they will likely move here."

Hoffman does share some of Steep's worry that Silicon Valley companies aren't paying enough attention to what the rest of the world is working on. He gets down sometimes about the endless "me-too-ism" of the business plans he sees—what he characterizes as "Airbnb for dog hotels." "If I could have a magic wand," he says, "I'd wish for more entrepreneurs to take greater and more diverse risks. And if those new ideas fit with the Valley model, then they should be here. And if not, they should think about building their company somewhere else."

Michael S. Malone has covered Silicon Valley for more than 25 years. His most recent book is The Intel Trinity.

21 Years Ago

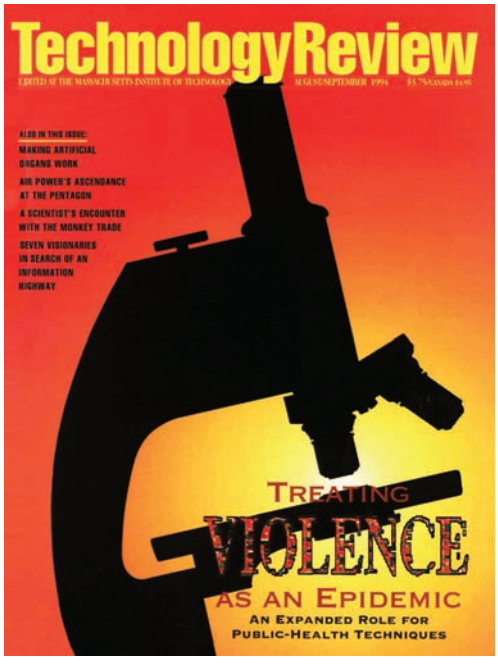

The Internet as a Force for Equality

Early in the Web era, a prominent technologist said the Net offered a way for the poor and the marginalized to gain power.

"Any vision of the information superhighway that focuses on video-on-demand and home shopping ... is too narrow. The first-order issue ought to be: What are we shooting for as a society? How are we conceiving of this great project that we are engaged in? My hope is that we reach a consensus for the system to be open, inclusive, egalitarian, and decentralized, and that it be based in the private sector so that investments can be matched with the possibility of reward.

One of the things that's gotten lost here is the notion of who we are as Americans. Our democratic tradition emphasizes the value of active participation in the shaping of one's society. If we have an information infrastructure that is highly open and decentralized and egalitarian and supports diversity, and that lets lots of people make lots of money, then it will create numerous opportunities for types of civic participation that do not exist today. So the point is to tip the balance back in favor of those who do not have lots of money and lots of power by giving them more of an opportunity to have their voices heard—by each other, by their elected officials, by people in their community who may not share their views.

Let's put this in historical perspective. In the first decade of the republic—when the United States was a start-up—it was Alexander Hamilton versus Thomas Jefferson for the dominant vision of what kind of government we were to have. Hamilton won, and we wound up with a highly centralized society. But now there's an opportunity to have sort of a rematch, under very different conditions, between the principles of Jefferson and those of Hamilton. We have the ability, given the construction of a high-capacity information infrastructure, to do things in a decentralized fashion that does not require large institutions, either public or private."

Excerpted from "Seven Thinkers in Search of the Information Highway," from the September 1994 issue of Technology Review. *The text above comes from Mitchell Kapor, chairman and cofounder of the Electronic Frontier Foundation and founder of Lotus Development.*

MIT Technology Review (ISSN 1099-274X), March/April 2015 Issue, Reg. U.S. Patent Office, is published bimonthly by MIT Technology Review, 1 Main St. Suite 13, Cambridge, MA 02142-1517. Periodicals postage paid at Boston, MA, and additional mailing offices. Postmaster: send address changes to MIT Technology Review, Subscriber Services, PO Box 5001, Big Sandy, TX 75755, or via the Internet at www.technologyreview.com/customerservice. Basic subscription rates: $41.94 per year within the United States; in all other countries, US$54. Publication Mail Agreement Number 40621028. Send undeliverable Canadian copies to PO Box 1051 Fort Erie, ON L2A 6C7. Printed in U.S.A. Audited by the Alliance for Audited Media

neat

WORK SMARTER, NOT HARDER

SAVE TIME WITH A
NEAT SCANNER + SMART ORGANIZATION SYSTEM.

Eliminate data entry and paperwork, streamline receipt and invoice management, and send information straight to QuickBooks and other applications – ***with Neat***.

neat.com/TR
866-395-2372

Get yours today at:

BEST BUY

Office DEPOT
OfficeMax

STAPLES
MAKE more HAPPEN

THE VALUE OF PERFORMANCE.

NORTHROP GRUMMAN

STEALTH

Not all stealth is created equal. For more than 70 years, Northrop Grumman has been the leader in tailless, blended-wing aircraft—an essential hallmark of true, low-observable stealth. And to this day, we're the only ones who have delivered large-scale, long-range stealth aircraft. ***THIS IS WHAT WE DO.***

www.northropgrumman.com/**stealth**